

Table of Contents

Table of Contents	1
F5 TMSH Reference - 16.x	17
General	17
grep	17
time	18
tmsh	19
Commands	23
cd	23
cp	24
create	25
delete	25
edit	26
exit	27
generate	28
help	28
install	29
list	29
load	31
modify	31
mv	32
publish	33
pwd	33
quit	34
reboot	34
reset-stats	35
restart	36
run	36
save	39
send-mail	40
show	40
shutdown	42
start	43
stop	43
submit	44
Modules	44
analytics	44
analytics afm-sweeper report	44
analytics afm-sweeper scheduled-report	47
analytics application-security-anomalies report	48
analytics application-security-anomalies scheduled-report	50
analytics application-security-incidents report	52
analytics application-security-network report	54
analytics application-security-network scheduled-report	57
analytics application-security report	58
analytics application-security scheduled-report	61
analytics asm-bypass report	63
analytics asm-bypass scheduled-report	65
analytics asm-cpu report	66
analytics asm-cpu scheduled-report	69
analytics asm-enforced-entities report	70
analytics asm-learning-suggestions report	72
analytics asm-memory report	75
analytics asm-memory scheduled-report	77
analytics asm-policy-changes report	78
analytics asm-violation report	81
analytics asm-violation scheduled-report	83
analytics bot-defense-event report	85
analytics cpu-per-vip report	87
analytics cpu report	89
analytics cpu scheduled-report	92
analytics device-traffic report	93
analytics device-traffic scheduled-report	96
analytics disk-info report	97
analytics disk-info scheduled-report	100

analytics dns-cache-resolver report	101
analytics dns-profile report	104
analytics dns-protocol scheduled-report	107
analytics dns-rpz report	109
analytics dns report	111
analytics dns scheduled-report	114
analytics dos-l3 report	115
analytics dos-l3 scheduled-report	118
analytics dos-l7 report	120
analytics dos-vis-attacks report	124
analytics dos-vis-common report	127
analytics dos-vis-vips report	130
analytics fw-nat report	133
analytics fw-nat scheduled-report	136
analytics global-settings	138
analytics gtm-wideip report	139
analytics http report	141
analytics http scheduled-report	146
analytics ip-intelligence report	148
analytics ip-intelligence scheduled-report	150
analytics ip-layer report	152
analytics ip-layer scheduled-report	154
analytics lsn-pool report	156
analytics lsn-pool scheduled-report	159
analytics memory-per-process report	160
analytics memory report	162
analytics memory scheduled-report	165
analytics network report	166
analytics network scheduled-report	170
analytics network stale-rules	172
analytics pem report	173
analytics pem scheduled-report	175
analytics pool-traffic report	177
analytics pool-traffic scheduled-report	179
analytics proc-cpu report	181
analytics proc-cpu scheduled-report	183
analytics protocol-inspection report	185
analytics protocol-security-http report	188
analytics protocol-security-http scheduled-report	190
analytics protocol-security report	192
analytics protocol-security scheduled-report	194
analytics report	195
analytics sip-dos report	200
analytics sip-dos scheduled-report	203
analytics sip report	205
analytics sip scheduled-report	207
analytics ssl-orchestrator-service-virtual report	209
analytics ssl-orchestrator-service-virtual scheduled-report	211
analytics ssl-orchestrator report	213
analytics ssl-orchestrator scheduled-report	216
analytics swg-blocked report	217
analytics swg-blocked scheduled-report	220
analytics swg report	221
analytics swg scheduled-report	223
analytics system-monitor report	225
analytics tcp-analytics report	228
analytics tcp-analytics scheduled-report	231
analytics tcp report	233
analytics tcp scheduled-report	235
analytics tmm-dns-zone report	237
analytics traffic-classification report	240
analytics traffic-classification scheduled-report	242
analytics udp report	244
analytics udp scheduled-report	247
analytics uri-type	248
analytics vcmp report	249
analytics vcmp scheduled-report	253
analytics virtual report	255
analytics virtual scheduled-report	258
api-protection	259
api-protection profile apiprotection	259
api-protection response	262

api-protection server	264
apm	265
apm aaa active-directory-trusted-domains	265
apm aaa active-directory	266
apm aaa crldp	268
apm aaa endpoint-management-system	269
apm aaa f5-mfa-configuration	271
apm aaa f5-service-connector	272
apm aaa http-connector-request	273
apm aaa http-connector-transport	275
apm aaa http	276
apm aaa kerberos-keytab-file	278
apm aaa kerberos	279
apm aaa ldap	280
apm aaa oam	282
apm aaa oauth-provider	284
apm aaa oauth-request	286
apm aaa oauth-server	287
apm aaa ocsp	289
apm aaa okta-connector	291
apm aaa radius	292
apm aaa saml-idp-automation	294
apm aaa saml-idp-connector	295
apm aaa saml	298
apm aaa securid	301
apm aaa tacacsplus	302
apm access-info	303
apm acl	305
apm apm-avr-config	307
apm client image	308
apm configuration captcha	309
apm epsec epsec-package	311
apm epsec software-status	312
apm license	312
apm log-setting	313
apm ntlm machine-account	315
apm ntlm ntlm-auth	316
apm oauth db-instance	317
apm oauth jwk-config	318
apm oauth jwt-config	320
apm oauth jwt-provider-list	322
apm oauth oauth-claim	323
apm oauth oauth-client-app	324
apm oauth oauth-resource-server	327
apm oauth oauth-scope	329
apm oauth purged-entries	330
apm oauth token-details	331
apm policy access-policy	332
apm policy agent aaa-active-directory	333
apm policy agent aaa-client-cert	335
apm policy agent aaa-crldp	336
apm policy agent aaa-http	337
apm policy agent aaa-ldap	338
apm policy agent aaa-oauth	340
apm policy agent aaa-ocsp	341
apm policy agent aaa-radius	342
apm policy agent aaa-saml	343
apm policy agent aaa-securid	344
apm policy agent acct-radius	346
apm policy agent acct-tacacsplus	347
apm policy agent api-authentication	348
apm policy agent api-server-selection	349
apm policy agent decision-box	350
apm policy agent dynamic-acl	351
apm policy agent ending-allow	352
apm policy agent ending-deny	352
apm policy agent ending-redirect	353
apm policy agent endpoint-check-machine-cert	355
apm policy agent endpoint-check-software	356
apm policy agent endpoint-linux-check-file	358
apm policy agent endpoint-linux-check-process	360
apm policy agent endpoint-mac-check-file	361

apm policy agent endpoint-mac-check-process	363
apm policy agent endpoint-machine-info	364
apm policy agent endpoint-windows-browser-cache-cleaner	365
apm policy agent endpoint-windows-check-file	366
apm policy agent endpoint-windows-check-process	368
apm policy agent endpoint-windows-check-registry	369
apm policy agent endpoint-windows-group-policy	370
apm policy agent endpoint-windows-info-os	372
apm policy agent endpoint-windows-protected-workspace	373
apm policy agent external-logon-page	374
apm policy agent http-header-modify	375
apm policy agent ip-geolocation-lookup	377
apm policy agent ip-reputation-lookup	378
apm policy agent irule-event	379
apm policy agent kerberos	380
apm policy agent l7-protocol-lookup	381
apm policy agent logging	382
apm policy agent logon-page	383
apm policy agent message-box	385
apm policy agent oam	386
apm policy agent oauth-authz	387
apm policy agent request-classification	391
apm policy agent resource-assign	392
apm policy agent response-selection	393
apm policy agent route-domain-selection	394
apm policy agent server-cert-response-control	395
apm policy agent server-cert-status	396
apm policy agent session-check	397
apm policy agent ssl-check	398
apm policy agent tacacsplus	399
apm policy agent variable-assign	400
apm policy customization-group	401
apm policy customization-languages	401
apm policy image-file	402
apm policy policy-item	402
apm policy windows-group-policy-file	402
apm profile access	403
apm profile connectivity	408
apm profile exchange	412
apm profile oauth	414
apm profile remote-desktop	418
apm profile vdi	419
apm report custom-report-field	420
apm resource app-tunnel	421
apm resource client-rate-class	423
apm resource client-traffic-classifier	424
apm resource ipv6-leasepool	426
apm resource leasepool	427
apm resource network-access	428
apm resource portal-access	433
apm resource remote-desktop citrix-client-bundle	435
apm resource remote-desktop citrix-client-package-file	436
apm resource remote-desktop citrix	437
apm resource remote-desktop quest	439
apm resource remote-desktop rdp	441
apm resource remote-desktop vmware-view	442
apm resource sandbox	444
apm resource webtop-link	445
apm resource webtop	446
apm saml artifact-resolution-service	448
apm saml attribute-consuming-service	449
apm saml auth-context-class-list	451
apm session	452
apm sso basic	452
apm sso form-based	454
apm sso form-basedv2	456
apm sso kerberos	461
apm sso ntlmv1	463
apm sso ntlmv2	465
apm sso oauth-bearer	466
apm sso saml-resource	468
apm sso saml-sp-automation	469

apm sso saml-sp-connector	470
apm sso saml	472
apm swg-content-type	475
apm swg-scheme	476
apm url-filter	476
asm	477
asm device-sync	478
asm http-method	478
asm httpclass-asm	479
asm policy	480
asm predefined-policy	482
asm response-code	483
asm webapp-language	484
auth	484
auth apm-auth	484
auth cert-ldap	485
auth ldap	489
auth login-failures	491
auth partition	492
auth password-policy	493
auth password	495
auth radius-server	495
auth radius	497
auth remote-role	498
auth remote-user	501
auth source	502
auth tacacs	503
auth user	504
cli	506
cli admin-partitions	506
cli alias private	507
cli alias shared	508
cli global-settings	510
cli history	511
cli preference	512
cli script	515
cli transaction	526
cli version	527
cm	528
cm add-to-trust	528
cm cert	529
cm config-sync	531
cm device-group	532
cm device	534
cm failover-status	537
cm key	538
cm remove-from-trust	539
cm sha1-fingerprint	540
cm sniff-updates	541
cm sync-status	541
cm traffic-group	542
cm trust-domain	544
cm watch-devicegroup-device	546
cm watch-sys-device	548
cm watch-trafficgroup-device	549
gtm	550
gtm datacenter	550
gtm distributed-app	551
gtm global-settings general	553
gtm global-settings load-balancing	557
gtm global-settings metrics-exclusions	558
gtm global-settings metrics	559
gtm iquery	561
gtm ldns	561
gtm link	562
gtm listener	565
gtm monitor bigip-link	568
gtm monitor bigip	569
gtm monitor external	571
gtm monitor firepass	573
gtm monitor ftp	575
gtm monitor gateway-icmp	577

gtm monitor gtp	579
gtm monitor http	581
gtm monitor https	583
gtm monitor imap	585
gtm monitor ldap	587
gtm monitor mssql	589
gtm monitor mysql	592
gtm monitor nntp	594
gtm monitor none	596
gtm monitor oracle	597
gtm monitor pop3	599
gtm monitor postgresql	601
gtm monitor radius-accounting	603
gtm monitor radius	605
gtm monitor real-server	607
gtm monitor scripted	609
gtm monitor sip	611
gtm monitor smtp	613
gtm monitor snmp-link	615
gtm monitor snmp	617
gtm monitor soap	619
gtm monitor tcp-half-open	621
gtm monitor tcp	623
gtm monitor udp	625
gtm monitor wap	627
gtm monitor wmi	629
gtm path	631
gtm persist	632
gtm pool a	633
gtm pool aaaa	639
gtm pool cname	646
gtm pool mx	652
gtm pool naptr	657
gtm pool srv	663
gtm prober-pool	669
gtm region	671
gtm rule	673
gtm server	674
gtm topology	679
gtm traffic	681
gtm wideip a	682
gtm wideip aaaa	685
gtm wideip cname	688
gtm wideip mx	690
gtm wideip naptr	693
gtm wideip srv	696
ilx	699
ilx global-settings	699
ilx node-version	700
ilx plugin	700
ilx workspace	709
ltm	714
ltm alg-log-profile	714
ltm auth crldp-server	716
ltm auth kerberos-delegation	717
ltm auth ldap	719
ltm auth ojsp-responder	721
ltm auth profile	724
ltm auth radius-server	726
ltm auth radius	727
ltm auth ssl-cc-ldap	729
ltm auth ssl-crldp	731
ltm auth ssl-ocsp	733
ltm auth tacacs	734
ltm cipher group	736
ltm cipher rule	737
ltm classification application	738
ltm classification auto-update settings	739
ltm classification auto-update status	739
ltm classification category	740
ltm classification ce	741
ltm classification signature-definition	742

ltm classification signature-update-schedule	743
ltm classification signature-version	744
ltm classification signatures	745
ltm classification stats application	746
ltm classification stats url-category	747
ltm classification stats urlcat-cloud	748
ltm classification update-signatures	749
ltm classification updates	749
ltm classification url-cat-policy	750
ltm classification url-category	751
ltm classification urldb-feed-list	752
ltm classification urldb-file	754
ltm clientssl-proxy cached-certs	755
ltm clientssl ocsdp-stapling-responses	755
ltm data-group external	756
ltm data-group internal	758
ltm default-node-monitor	759
ltm dns analytics global-settings	760
ltm dns cache global-settings	761
ltm dns cache records all	762
ltm dns cache records key	762
ltm dns cache records msg	763
ltm dns cache records nameserver	764
ltm dns cache records rrset	765
ltm dns cache resolver	766
ltm dns cache transparent	770
ltm dns cache validating-resolver	772
ltm dns dns-express-db	775
ltm dns dnssec key	776
ltm dns dnssec zone	778
ltm dns nameserver	780
ltm dns tsig-key	782
ltm dns zone	783
ltm eviction-policy	784
ltm global-settings connection	788
ltm global-settings general	789
ltm global-settings rule	790
ltm global-settings traffic-control	791
ltm ifile	793
ltm lsn-log-profile	794
ltm lsn-pool	795
ltm message-routing diameter peer	799
ltm message-routing diameter profile router	801
ltm message-routing diameter profile session	803
ltm message-routing diameter route	808
ltm message-routing diameter transport-config	809
ltm message-routing generic peer	810
ltm message-routing generic protocol	812
ltm message-routing generic route	814
ltm message-routing generic router	815
ltm message-routing generic transport-config	817
ltm message-routing mqtt peer	818
ltm message-routing mqtt profile router	819
ltm message-routing mqtt profile session	821
ltm message-routing mqtt route	822
ltm message-routing mqtt transport-config	824
ltm message-routing sip peer	825
ltm message-routing sip profile router	827
ltm message-routing sip profile session	830
ltm message-routing sip route	833
ltm message-routing sip transport-config	834
ltm monitor diameter	836
ltm monitor dns	838
ltm monitor external	842
ltm monitor firepass	844
ltm monitor ftp	847
ltm monitor gateway-icmp	849
ltm monitor http	852
ltm monitor http2	855
ltm monitor https	859
ltm monitor icmp	862
ltm monitor imap	865

ltm monitor inband	868
ltm monitor ldap	869
ltm monitor module-score	872
ltm monitor mqtt	874
ltm monitor mssql	876
ltm monitor mysql	879
ltm monitor nntp	882
ltm monitor none	885
ltm monitor oracle	886
ltm monitor pop3	888
ltm monitor postgresql	891
ltm monitor radius-accounting	894
ltm monitor radius	896
ltm monitor real-server	899
ltm monitor rpc	900
ltm monitor sasp	903
ltm monitor scripted	904
ltm monitor sip	907
ltm monitor smb	910
ltm monitor smtp	913
ltm monitor snmp-dca-base	915
ltm monitor snmp-dca	917
ltm monitor soap	919
ltm monitor tcp-echo	922
ltm monitor tcp-half-open	925
ltm monitor tcp	927
ltm monitor udp	930
ltm monitor virtual-location	933
ltm monitor wap	935
ltm monitor wmi	938
ltm nat-stats	940
ltm nat	941
ltm node	943
ltm persistence cookie	945
ltm persistence dest-addr	948
ltm persistence global-settings	950
ltm persistence hash	951
ltm persistence host	953
ltm persistence msrdp	955
ltm persistence persist-records	956
ltm persistence sip	958
ltm persistence source-addr	960
ltm persistence ssl	962
ltm persistence universal	964
ltm policy-strategy	965
ltm policy	968
ltm pool	988
ltm profile analytics	995
ltm profile certificate-authority	1001
ltm profile classification	1003
ltm profile client-ldap	1004
ltm profile client-ssl	1005
ltm profile dhcpv4	1012
ltm profile dhcpv6	1016
ltm profile diameter	1019
ltm profile dns-logging	1022
ltm profile dns	1024
ltm profile fasthttp	1026
ltm profile fastl4	1029
ltm profile fix	1034
ltm profile ftp	1036
ltm profile georedundancy	1038
ltm profile gtp	1039
ltm profile html	1041
ltm profile http-compression	1042
ltm profile http	1045
ltm profile http2	1051
ltm profile http3	1053
ltm profile httprouter	1054
ltm profile icap	1055
ltm profile iiop	1056
ltm profile ilx	1058

ltm profile imap	1059
ltm profile ipother	1061
ltm profile ipsecalg	1062
ltm profile mapt	1063
ltm profile mblb	1065
ltm profile mqtt	1067
ltm profile mssql	1068
ltm profile netflow	1069
ltm profile ntlm	1071
ltm profile oosp-stapling-params	1072
ltm profile oosp	1073
ltm profile one-connect	1074
ltm profile pcp	1076
ltm profile pop3	1078
ltm profile pptp	1079
ltm profile qoe	1080
ltm profile quic	1081
ltm profile radius	1083
ltm profile ramcache	1084
ltm profile request-adapt	1085
ltm profile request-log	1087
ltm profile response-adapt	1089
ltm profile rewrite	1091
ltm profile rtsp	1094
ltm profile sctp	1096
ltm profile server-ldap	1099
ltm profile server-ssl	1100
ltm profile sip	1106
ltm profile smtp	1109
ltm profile smtps	1110
ltm profile socks	1111
ltm profile splitsessionclient	1112
ltm profile splitsessionserver	1114
ltm profile statistics	1115
ltm profile stream	1117
ltm profile tcp-analytics	1119
ltm profile tcp	1120
ltm profile tftp	1127
ltm profile traffic-acceleration	1128
ltm profile udp	1130
ltm profile wa-cache	1132
ltm profile web-acceleration	1132
ltm profile web-security	1134
ltm profile websocket	1135
ltm profile xml	1137
ltm rule-profiler	1138
ltm rule	1140
ltm snat-translation	1142
ltm snat	1144
ltm snatpool	1146
ltm tacdb customdb-file	1147
ltm tacdb customdb	1148
ltm tacdb licenseddb	1149
ltm tacdb query	1150
ltm traffic-class	1151
ltm traffic-matching-criteria	1152
ltm urlcat-cloud-cache	1154
ltm urlcat-query	1155
ltm virtual-address	1155
ltm virtual	1158
mgmt	1164
mgmt shared settings api-status availability	1165
mgmt shared settings api-status log resource-property	1166
mgmt shared settings api-status log resource	1167
net	1168
net address-list	1168
net arp	1170
net bwc policy	1171
net bwc priority-group	1177
net bwc traffic-group	1178
net clone-stats	1180
net cmetrics	1180

net cos global-settings	1181
net cos map-8021p	1182
net cos map-dscp	1183
net cos traffic-priority	1184
net dag-globals	1185
net dns-resolver	1186
net f5optics	1188
net fdb tunnel	1189
net fdb vlan	1190
net ike-evt-stat	1192
net ike-msg-stat	1192
net interface-cos	1193
net interface-ddm	1193
net interface	1194
net ipsec-stat	1198
net ipsec ike-daemon	1198
net ipsec ike-peer	1199
net ipsec ike-sa	1202
net ipsec ipsec-policy	1203
net ipsec ipsec-sa	1205
net ipsec manual-security-association	1205
net ipsec traffic-selector	1207
net ipv6-subscriber-prefix-length	1208
net lldp-globals	1209
net lldp-neighbors	1209
net mroute	1210
net multicast-globals	1211
net ndp	1212
net packet-filter-trusted	1213
net packet-filter	1214
net packet-tester security	1217
net port-list	1219
net port-mirror	1220
net rate-shaping class	1221
net rate-shaping color-policer	1223
net rate-shaping drop-policy	1225
net rate-shaping queue	1226
net rate-shaping shaping-policy	1228
net route-domain	1230
net route	1232
net router-advertisement	1234
net routing access-list	1236
net routing bfd	1237
net routing bgp	1238
net routing community-list	1247
net routing debug	1248
net routing extcommunity-list	1248
net routing prefix-list	1249
net routing profile bgp	1249
net routing route-map	1251
net rst-cause	1254
net self-allow	1254
net self	1255
net service-policy	1258
net sfc-stats	1259
net sfc chain	1260
net sfc hop	1261
net sfc sf	1262
net stp-globals	1263
net stp	1265
net timer-policy	1266
net trunk	1269
net tunnels endpoint	1272
net tunnels etherip	1273
net tunnels fec-stat	1274
net tunnels fec	1275
net tunnels geneve	1276
net tunnels gre	1278
net tunnels ipip	1279
net tunnels ipsec	1280
net tunnels lw4o6	1281
net tunnels map	1283

net tunnels ppp	1284
net tunnels tcp-forward	1285
net tunnels tunnel	1286
net tunnels v6rd	1288
net tunnels vxlan	1290
net tunnels wccp	1291
net vlan-allowed	1292
net vlan-group	1293
net vlan	1295
net wccp	1298
pem	1300
pem forwarding-endpoint	1300
pem global-settings analytics	1303
pem global-settings gx	1304
pem global-settings hsl-flow	1305
pem global-settings hsl-report	1305
pem global-settings insert-content	1306
pem global-settings policy	1307
pem global-settings quota-mgmt	1308
pem global-settings session-mgmt-attributes	1309
pem global-settings subscriber-activity-log	1310
pem interception-endpoint	1312
pem irule	1313
pem listener	1314
pem policy	1316
pem profile diameter-endpoint	1327
pem profile radius-aaa	1329
pem profile spm	1330
pem profile subscriber-mgmt	1332
pem protocol diameter-avp	1333
pem protocol profile gx	1335
pem protocol profile radius	1338
pem protocol radius-avp	1340
pem quota-mgmt rating-group	1342
pem reporting format-script	1344
pem service-chain-endpoint	1347
pem sessiondb	1349
pem stats action	1353
pem stats dtos	1354
pem stats gx	1355
pem stats gy	1357
pem stats hsl	1358
pem stats hudnode-opt	1359
pem stats multiple-ip	1359
pem stats persistence	1360
pem stats radius	1361
pem stats sd	1362
pem stats subscriber	1363
pem stats tethering	1366
pem subscriber-attribute	1367
pem subscriber	1369
pem subscribers	1370
security	1371
security analytics settings	1371
security anti-fraud engine-update	1373
security anti-fraud profile	1374
security anti-fraud signatures-update	1399
security blacklist-publisher all-blacklist-publisher	1401
security blacklist-publisher blacklist-publisher-stats	1401
security blacklist-publisher by-addr	1402
security blacklist-publisher by-category	1403
security blacklist-publisher category	1404
security blacklist-publisher profile	1405
security bot-defense anomaly-category	1406
security bot-defense anomaly	1407
security bot-defense class	1407
security bot-defense micro-service	1408
security bot-defense profile	1408
security bot-defense signature-category	1417
security bot-defense signature	1418
security bot-defense template	1419
security cloud-services cmd	1420

security cloud-services connector	1421
security datasync background-tasks	1422
security datasync device-stats	1423
security datasync global-profile	1424
security datasync local-profile	1425
security debug drop-redirect-stats	1427
security debug matcher	1427
security debug register	1428
security device-id attribute	1430
security device device-context	1431
security dos auto-thresholds heavy-urls	1431
security dos auto-thresholds stress-based	1432
security dos auto-thresholds top-device-ids	1433
security dos auto-thresholds top-geolocations	1433
security dos auto-thresholds top-source-ips	1434
security dos auto-thresholds top-urls	1434
security dos auto-thresholds tps-based	1435
security dos autodos-file-object	1435
security dos behavioral-signature	1436
security dos bot-signature-category	1437
security dos bot-signature	1438
security dos device-config	1439
security dos dns-nxdomain-stat	1448
security dos dos-signature	1449
security dos dynamic-signatures	1452
security dos ip-uncommon-protolist	1453
security dos l4bdos-file-object	1454
security dos network-whitelist	1455
security dos profile	1459
security dos spva-stats	1473
security dos stress-stats	1475
security dos udp-portlist	1476
security dos virtual	1478
security firewall address-list	1479
security firewall config-change-log	1481
security firewall container-stat	1481
security firewall context-stat	1482
security firewall current-state	1482
security firewall fqdn-entity	1483
security firewall fqdn-info	1484
security firewall global-fqdn-policy	1484
security firewall global-rules	1485
security firewall ipi-category-info	1487
security firewall management-ip-rules	1487
security firewall matching-rule	1492
security firewall on-demand-compilation	1492
security firewall on-demand-rule-deploy	1493
security firewall policy	1493
security firewall port-list	1497
security firewall port-misuse-policy	1499
security firewall rule-list	1501
security firewall rule-stat	1507
security firewall schedule	1508
security firewall user-domain	1509
security firewall user-list	1510
security firewall uuid-default-autogenerate	1511
security flowspec-route-injector flowspec-advertised-route-info	1512
security flowspec-route-injector profile	1512
security http file-type	1516
security http mandatory-header	1517
security http profile	1518
security ip-intelligence blacklist-category	1522
security ip-intelligence feed-list	1523
security ip-intelligence global-policy	1525
security ip-intelligence info	1526
security ip-intelligence policy	1527
security log antifraud-storage-field	1529
security log network-storage-field	1530
security log profile	1530
security log protocol-dns-storage-field	1543
security log protocol-sip-storage-field	1543
security log remote-format	1544

security log storage-field	1545
security malicious-sources device-ids	1546
security malicious-sources ip-addresses	1546
security nat destination-translation	1547
security nat policy	1548
security nat source-translation	1551
security packet-filter default-rules	1556
security packet-filter policy	1556
security packet-filter rule-stat	1558
security presentation tmui netflow-details	1558
security presentation tmui netflow-list	1559
security presentation tmui signature-details	1559
security presentation tmui signature-list	1560
security protected-servers netflow-tmc-stat	1560
security protocol-inspection auto-update settings	1561
security protocol-inspection auto-update status	1561
security protocol-inspection common-config	1562
security protocol-inspection compliance-enums	1563
security protocol-inspection compliance	1563
security protocol-inspection learning-stats	1565
security protocol-inspection learning-suggestions	1565
security protocol-inspection profile-status	1566
security protocol-inspection profile	1566
security protocol-inspection service	1568
security protocol-inspection signature	1569
security protocol-inspection staging	1570
security protocol-inspection system	1571
security protocol-inspection updates	1572
security protocol-inspection virtual-servers	1572
security scrubber dwbl-scrubber-category-stats	1573
security scrubber dwbl-scrubber-stat	1573
security scrubber profile	1574
security scrubber unredirect	1578
security ssh profile	1579
security zone	1582
sys	1583
sys air-filter-reset	1583
sys alert lcd	1584
sys aom	1584
sys appiq config	1585
sys application apl-script	1586
sys application custom-stat	1587
sys application service	1588
sys application template	1590
sys autoscale-group	1598
sys availability	1598
sys clock	1599
sys cluster	1600
sys config-diff	1601
sys config	1602
sys connection	1606
sys console	1607
sys core	1608
sys cpu	1609
sys crypto acceleration-strategy	1610
sys crypto allow-key-export	1610
sys crypto ca-bundle-manager	1611
sys crypto cert-order-manager	1612
sys crypto cert-validation-response ocsp	1614
sys crypto cert-validator crl	1614
sys crypto cert-validator ocsp	1615
sys crypto cert	1617
sys crypto check-cert	1619
sys crypto client	1620
sys crypto crl	1621
sys crypto csr	1622
sys crypto encrypted-attributes	1624
sys crypto fips by-handle	1624
sys crypto fips external-hsm	1625
sys crypto fips key	1625
sys crypto key	1626
sys crypto master-key	1630

sys crypto pkcs12	1630
sys crypto server	1632
sys daemon-ha	1633
sys daemon-log-settings clusterd	1634
sys daemon-log-settings csyncd	1635
sys daemon-log-settings icr-eventd	1635
sys daemon-log-settings icrd	1636
sys daemon-log-settings lind	1637
sys daemon-log-settings mcpd	1638
sys daemon-log-settings tmm	1638
sys datastor	1640
sys db	1641
sys default-config	1642
sys diags ihealth-request	1643
sys diags ihealth-result	1644
sys diags ihealth	1645
sys disk application-volume	1646
sys disk directory	1646
sys disk logical-disk	1647
sys dns	1648
sys dynad instrumentation	1649
sys dynad key	1650
sys dynad rpm	1651
sys dynad settings	1652
sys dynad status	1653
sys ecm config	1653
sys ecm register	1654
sys failover	1654
sys feature-module	1656
sys file apache-ssl-cert	1657
sys file browser-capabilities-db	1658
sys file data-group	1659
sys file device-capabilities-db	1661
sys file external-monitor	1662
sys file ifile	1663
sys file lwtunneltbl	1664
sys file rewrite-rule	1665
sys file ssl-cert	1666
sys file ssl-crl	1668
sys file ssl-key	1669
sys fipsuser	1671
sys fix-connection	1671
sys folder	1672
sys fpga firmware-config	1673
sys fpga info	1674
sys fpga turboflex-profile	1675
sys geoip	1675
sys global-settings	1676
sys ha-group	1679
sys ha-status	1680
sys hardware	1681
sys host-info	1682
sys httpd	1682
sys hypervisor-info	1685
sys icall event	1686
sys icall handler periodic	1687
sys icall handler perpetual	1688
sys icall handler triggered	1689
sys icall istats-trigger	1691
sys icall publisher	1692
sys icall script	1692
sys icmp-stat	1694
sys icontrol-soap	1695
sys integrity status-check	1696
sys internal-proxy	1696
sys ip-address	1697
sys ip-stat	1698
sys ipfix destination	1698
sys ipfix element	1699
sys ipfix irules	1700
sys iprep-status	1701
sys license	1702

sys log-config destination alertd	1703
sys log-config destination arcsight	1704
sys log-config destination ipfix	1705
sys log-config destination local-database	1706
sys log-config destination local-syslog	1707
sys log-config destination management-port	1708
sys log-config destination remote-high-speed-log	1709
sys log-config destination remote-syslog	1711
sys log-config destination splunk	1712
sys log-config filter	1713
sys log-config publisher	1715
sys log-rotate	1716
sys log	1718
sys mac-address	1719
sys management-dhcp	1720
sys management-ip	1721
sys management-ovsdb	1722
sys management-proxy-config	1723
sys management-route	1724
sys mcp-state	1726
sys memory	1726
sys nethsm async-queue-stat	1727
sys nethsm pkcs11d-stat	1728
sys nethsm sync-queue-stat	1728
sys ntp	1729
sys outbound-smtp	1731
sys performance all-stats	1732
sys performance connections	1732
sys performance dnsexpress	1733
sys performance dnssec	1734
sys performance gtm	1734
sys performance ramcache	1735
sys performance system	1735
sys performance throughput	1736
sys pfman consumer	1737
sys pfman device	1738
sys proc-info	1738
sys provision	1739
sys pva-traffic	1741
sys raid array	1741
sys raid bay	1742
sys raid disk	1743
sys ready	1744
sys scriptd	1744
sys service	1745
sys sflow data-source http	1746
sys sflow data-source interface	1747
sys sflow data-source system	1747
sys sflow data-source vlan	1748
sys sflow global-settings http	1748
sys sflow global-settings interface	1749
sys sflow global-settings system	1750
sys sflow global-settings vlan	1751
sys sflow receiver	1751
sys smtp-server	1753
sys snmp	1754
sys software block-device-hotfix	1759
sys software block-device-image	1761
sys software hotfix	1762
sys software image	1764
sys software signature	1766
sys software status	1767
sys software update-status	1767
sys software update	1769
sys software volume	1770
sys sshd	1771
sys state-mirroring	1773
sys sync-sys-files	1774
sys syslog	1775
sys tmm-info	1777
sys tmm-traffic	1778
sys traffic	1778

sys turboflex features	1779
sys turboflex profile-config	1780
sys turboflex profile all	1780
sys turboflex profile feature	1781
sys turboflex warning	1782
sys ucs	1782
sys url-db download-result	1784
sys url-db download-schedule	1784
sys url-db url-category	1785
sys version	1787
util	1787
util ccmode	1787
util clientssl-ciphers	1788
util diadb	1789
util dnatutil	1789
util establish adfs trust	1790
util finalize custom ami	1791
util geodb	1791
util geoutil	1792
util ihealth	1792
util ipsecalgdb	1793
util lsndb	1794
util platform check	1796
util platform diag	1797
util qkcloud	1797
util serverssl-ciphers	1798
util sipdb	1798
util ssh keyswap	1799
util test-monitor	1799
util verify encryption	1799
vcmp	1800
vcmp global	1800
vcmp guest	1801
vcmp health ha-status	1804
vcmp health module-provision	1804
vcmp health prompt	1805
vcmp health software	1806
vcmp traffic-profile	1806
vcmp virtual-disk-template	1807
vcmp virtual-disk	1808
wam	1808
wam ad-policy	1808
wam application	1809
wam domain list	1812
wam object-type	1813
wam policy	1815
wam resource concat-set	1830
wam resource domain-list	1831
wam resource url	1832
wam roi-statistics	1832
wom	1833
wom advertised-route	1833
wom deduplication	1834
wom diagnose-conn	1835
wom endpoint-discovery	1836
wom local-endpoint	1837
wom profile cifs	1839
wom profile isession	1841
wom profile mapi	1843
wom remote-endpoint	1844
wom remote-route	1846
wom server-discovery	1847
wom verify-config	1849

F5 TMSH Reference - 16.x



F5 TMSH references are collections of the available BIG-IP TMSH man pages.

General

grep

NAME

grep - Display lines matching a pattern

SYNTAX

```
list [component] "|" grep [ [option | pattern] ... ]
show [component] "|" grep [ [option | pattern] ... ]
```

options:

- A [integer]
- B [integer]
- C [integer]
- E
- G
- P
- c
- e [pattern]
- i
- m [integer]
- n
- o
- v
- w

-x Note: Each option must be followed by a space.

Note: tmsd treats any argument that is not preceded by a supported option, and does not begin with a hyphen, as a search pattern preceded by -e.

DESCRIPTION

You can use grep to filter the output generated by the commands list (configuration settings) and show (statistics and runtime status). You must type the character | before the grep specification. You can use multiple filters chained together.

EXAMPLES

The following examples show how to use the grep utility in tmsd.

```
list ltm node | grep "^10\."
list ltm virtual | grep -i seattle
list ltm virtual | grep -i abc | grep -i ab | grep -i a
```

OPTIONS

-A Display the specified number of lines of context after matching lines.

-B Display the specified number of lines of context before matching lines.

-C Display the specified number of lines of context before and after matching lines.

-E Interpret patterns as extended regular expressions.

-G Interpret patterns as basic regular expressions.

-P Interpret patterns as Perl regular expressions.

-c Display a count of the lines that match. If -v is specified the number of non-matching lines is displayed.

-e Specify a pattern. This is useful to protect against patterns beginning with a hyphen.

-i Case insensitive search.

-m Stop reading input after the specified number of matching lines. If -c is specified the count will not exceed the value specified for -m. If -v is specified grep will stop after finding the specified number of non-matching lines.

-n Prefix each line of output with the line number relative to the input.

-o Show only the part of a matching line that matches the pattern.

-v Invert the sense of matching, to select non-matching lines.

-w Select only those lines containing matches that form whole words. The test is that the matching substring must either be at the beginning of the line, or preceded by a non-word constituent character. Similarly, it must be either at the end of the line or followed by a non-word constituent character. Word-constituent characters are letters, digits, and the underscore.

-x Select only those matches that exactly match the whole line.

SEE ALSO

list, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009. All rights reserved.

BIG-IP 2017-05-24 grep(1)

time

NAME

Time - Date and Time formats.

MODULE

All tmsb modules.

SYNTAX

Date/Time Syntax

```
now[ [ + | - ] [ d | h | w | m ] ]
yyyy-mm-dd[ : | T ]hh:mm[:ss]
mm-dd[-yyyy][ : | T ]hh:mm[:ss]
mm/dd[/yyyy][ : | T ]hh:mm[:ss]
```

Date Range Syntax

```
now[ [ + | - ] [ d | h | w | m ] ]--now[ [ + | - ] [ d | h | w | m ] ]
yyyy-mm-dd[ : | T ]hh:mm[:ss]--yyyy-mm-dd[ : | T ]hh:mm[:ss]
mm-dd[-yyyy][ : | T ]hh:mm[:ss]-indefinite
epoch--mm/dd[/yyyy][ : | T ]hh:mm[:ss]
now[ [ + | - ] [ d | h | w | m ] ]
```

DESCRIPTION

The date or time format is found in tmsb as an attribute or parameter for many configuration items. Below are the various formats supported for both Date/Time and Date Range. Please see the examples for further assistance in using the required formats.

DATE:TIME FORMATS

nowX This date format starts with now (the current time) and is optionally followed by + or - some time span. The format will look like the following: now[[+ | -] integer [d | h | w | m]], where the user picks either before (-) or after (+) the current time and then specifies integer number of minutes(m), hours(h), days(d) or weeks(w). This format is case-insensitive.

Examples:

Input Date	Description
------------	-------------

now-3d	3 days ago.
now+3h	3 hours from now.
now-3m	3 minutes ago.
now+3w	3 weeks from now.

yyyy-mm-dd:hh:mm:ss

This format requires a year, month, day separated by - characters. A time is also required, which is specified as hour:minute:second, where the seconds are optional. The date and time must be separated by a : colon. Note: This is the default time format for output from tmsb.

Examples:

Input Date	Description
------------	-------------

2013-05-29:13:30	May 29th, 2013 at 1:30pm.
2000-01-04:12:22:30	January 4th, 2000 at 12:22pm and 30 seconds.

mm-dd-yyyy:hh:mm:ss

This format requires at least a month(m) and day(d) specified and optionally a year (y). If no year is specified, tmsb will auto-fill the year with the current year. A time is also required in the format of hour:minute:second, where the seconds are optional.

Examples:

Input Date	Description
------------	-------------

3-12-2015:12:01:00 March 12th, 2015 at 12:01 pm.
4-15:22:10:30 April 15th of this year at 10:10 pm and 30 seconds.

mm/dd/yyyy:hh:mm:ss

This format requires at least a month(m) and day(d) specified and optionally a year (y). If no year is specified, tmsh will auto-fill the year with the current year. A time is also required in the format of hour:minute:second, where the seconds are optional.

Examples:

Input Date	Description
------------	-------------

3/12/2015:12:01:00	March 12th, 2015 at 12:01 pm.
4/15:22:10:30	April 15th of this year at 10:10 pm and 30 seconds.

T Delimiter

Any of the above time formats may optionally use a capital letter T (as in the word Time) to separate the date from the time, instead of using a colon (:).

Examples:

Input Date	Description
------------	-------------

9/16/2005T12:01:01	September 16th, 2005 at 12:01pm and 1 second.
2011-11-12T00:03:30	November 12th, 2011 at 12:03am and 30 seconds.

Special Dates

There are two special dates that may be used in tmsh. They are indefinite and epoch. Below is an explanation of those dates.

indefinite

The date will be marked as being infinitely in the future (end of time).

epoch

The date will be marked as being infinitely in the past (beginning of time).

DATE RANGES

DateX--DateZ

A Date Range is 2 dates in a valid Date Format separated by a -- (double hyphen). The dates may be any of the Date Formats specified above. See examples below on how to use this notation.

Examples:

Input Date	Description
------------	-------------

now-2d--now-4d	2 to 4 days ago.
now--now-3m	From 3 minutes ago to now.
epoch--3/12/2011:12:00:00	Everything older than March 12th, 2011 at noon.
2008-03-12--indefinite	Everything after midnight on March 12th, 2008.

DateX

When specifying a date range, the second date may be left out. This will cause the system to assume the second date in the range to be now. Using this format for a date range may make it confusing when using the NowX date format listed above. The following examples will help clarify how to use this format with any supported Date Format.

Examples:

Input Date	Description
------------	-------------

now-3d	From 3 days ago to now.
now+3w	From now to 3 weeks from now.
epoch	Everything before the current date and time.
indefinite	Everything after the current date and time.

SEE ALSO

tmsh, create, modify

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 time(1)

tmsh

NAME

tmsh - Traffic Management Shell - A command line interface for managing the BIG-IP(r) system.

DESCRIPTION

You can use tmsh to configure and manage the BIG-IP system in conjunction with the Configuration utility, which is the browser-based BIG-IP system and network management tool.

MODULES

The structure of tmsh is hierarchical and modular. The highest level is the root module, which contains subordinate modules: auth, cli, gtm, ltm, net, sys and wom. Use the command help with no arguments to display the module hierarchy relative to the current module.

The gtm, ltm, net, sys, and wom modules also contain subordinate modules. All modules and subordinate modules contain components. To display the list of modules and components that are available in the current module type Tab or ? at the tmsh prompt.

Commands operate on components. To display the list of available commands type Tab or ? at the beginning of the command line. To display a list of components on which a command can operate type the command followed by a space followed by Tab or ?.

The following examples illustrate how to navigate the tmsh hierarchy.

To enter a module, type the name of the module at the tmsh prompt.

```
(tmsh)# ltm
```

The prompt displays the current module location.

```
(tmsh.ltm)#
```

You can display the components in a module using the commands list (configuration) and show (statistics and runtime status). The following command sequence displays the virtual server configuration of the BIG-IP system.

```
(tmsh.ltm)# list virtual
```

In the following examples, the commands list and show display information about only ltm components.

```
(tmsh.ltm)# list  
(tmsh.ltm)# show
```

You can access any component in any module from any other module by specifying a complete path to the component. For example, from the ltm module, the following command displays all of the properties of the VLANs on the system. The forward slash / specifies that what follows is relative to the root module.

```
(tmsh.ltm)# list /net vlan all-properties
```

The forward slash is optional if the root module is the current module. For example, the following command sequences display profiles.

```
(tmsh)# list ltm profile  
(tmsh)# list /ltm profile  
(tmsh)# list / ltm profile
```

Most components also support component mode. You can navigate to a single component and run commands to manage that component. For example, from the ltm module, to navigate to the node component, use the following command.

```
(tmsh.ltm)# node
```

To display the properties of all nodes, use the following command.

```
(tmsh.ltm.node)# list
```

You can also navigate to a specific object (object mode). For example, from the node component, to enter object mode for a specific node, enter the command modify followed by the IP address of the node.

```
(tmsh.ltm.node)# modify 10.1.1.10
```

In object mode, you can configure property settings directly. For example, to set the connection limit for 10.1.1.10 to 10000, use the following command.

```
(tmsh.ltm.node.10.1.1.10)# connection-limit 10000
```

To exit a module enter the command exit at the tmsh prompt, as shown below.

```
(tmsh.ltm)# exit  
(tmsh)#
```

PRODUCT PROVISIONING

You must provision a BIG-IP system module before you can use tmsh to configure that product, for example, the Global Traffic Manager. The command sequence list sys provision displays the BIG-IP system modules that can be provisioned. For more information about provisioning, see the TMOS(r) Management Guide for BIG-IP Systems and help sys provision.

LOADING/SAVING THE SYSTEM CONFIGURATION

The system applies all configuration changes that you make from within tmsh to the running configuration of the system.

You can save a portion of the running configuration known as the base configuration. You can also load the base configuration from the stored configuration files.

To save the base configuration to the stored configuration files, use the command sequence: `save sys base-config`.

To replace the running base configuration with the configuration in the stored configuration files, use the command sequence: `load /sys base-config`.

Additionally, you can save the entire running configuration or load all of the stored configuration files.

To save the entire running configuration to the stored configuration files, use the command sequence: `save /sys config`.

To replace the entire running configuration with the configuration in the stored configuration files using the command sequence: `load /sys config`.

HELP

tmsh includes man pages for each of the commands and components that are available within tmsh. You access the man pages using the following command syntax: `help [[command] | [full path to component]]`.

For example, to access the man page for the vlan component from the root module, use this command sequence: `help / net vlan`.

You can also search the man pages for information on a specific topic. To do this you use the command syntax: `help search [topic]`. You can perform a help search from within any module in the tmsh hierarchy. For example, to find the man pages that contain a reference to VLANs, use this command sequence: `help search vlan`

To display a list of topics that are available in a module use this command sequence: `help [full path to module]`.

For example, to display the topics that are available in the current module use this command: `help`. To display the topics that are available in the net module use this command sequence: `help / net`.

CONTEXT-SENSITIVE HELP

tmsh includes a context-sensitive help feature that provides help as you type commands. At any time, you can type a question mark (?) on the command line, and tmsh returns information to assist you in completing the command. Based on when you type the question mark, you get the following results.

When you type a question mark immediately following any portion of a command, tmsh returns possible completions for the command, but does not complete the command as the command completion feature does. When you type a space before the question mark, tmsh returns descriptive text that explains the commands, components, or properties that you can configure.

When you type a question mark in the middle of a command, tmsh returns help on the command to the left of the cursor.

Note: To use a question mark in a Glob or regular expression, you must escape the question mark using quotation marks, apostrophes, or a backslash.

Additionally, you can request context-sensitive help for the last command in a series of commands. For more information, see ENTERING MULTIPLE COMMANDS, following.

COMMAND COMPLETION

At any point while typing or editing a command in tmsh, you can press the Tab key. tmsh either completes the current or next word, or displays possible completions for the current or next word. If tmsh displays nothing after you press the Tab key, no options exist to complete the word. If you move the cursor anywhere on the command line and press the Tab key, tmsh completes what is to the left of the cursor.

Command completion also reduces the amount of typing that is required to run commands. When you press the Tab key, the system automatically completes the current command-line element to as many unique characters as possible. If there is more than one possible completion the list of possible completions displays. Command completion also completes configuration object identifiers.

ENTERING MULTIPLE COMMANDS

You can enter multiple commands on the command line by separating the commands with semi-colons (;). For example, to display the properties of the self IP addresses and VLANs of the system, use this command sequence:

```
list / net self ; list / net vlan
```

When you enter multiple commands in this way, all of the commands are added to the command history in a single line item, regardless of whether any of the commands were successful. However, if one of the commands that you enter fails to parse, tmsh does not run the remaining commands you entered. tmsh audits commands as the commands run; therefore, if a command fails to parse, tmsh does not audit the remaining commands. For more information about the command history, see COMMAND HISTORY, following.

You can also specify multiple commands in a command alias by separating the commands with semi-colons. For example, to create an alias that displays the properties of the VLANs and VLAN groups on the system, use this command sequence:

```
create / cli alias vlans command "list / net vlan ; list / net vlan-group"
```

You can request context-sensitive help and utilize the command completion feature on the last command in a series of commands. For example, the following command sequence displays help for the vlan-group component.

```
list / net vlan ; list / net vlan-group ?
```

COMMAND HISTORY

tmsh saves in the command history file each command that you enter. The command history persists when you log off of the system. The next time you log on to the system, you can search for, display, and then edit, the tmsh commands that you entered in previous sessions. The command history persists even through a restart of

the BIG-IP system. For more information about the command history feature, see help history.

The following examples show how to use the command history feature.

To display the commands in the history list, enter either the command sequence show history or an exclamation point (!). tmsh displays a list of commands each preceded by a numeric ID.

To run a command from the history list, enter an exclamation point followed by the numeric ID of the command.

To run the previous command, enter !!.

FILTERING OUTPUT

You can filter the output generated by the commands list (configuration settings) and show (statistics and runtime status) using the UNIX grep utility. You must type the character | before the grep specification. You can use multiple filters chained together. For a list of supported grep options, see the Traffic Management Shell (tmsh) Reference Guide.

The following examples show how to use the grep utility in tmsh.

```
list ltm node | grep "^10\.2"
list ltm virtual | grep -i seattle
list ltm virtual | grep -i abc | grep -i ab | grep -i a
```

KEYBOARD BINDINGS

tmsh supports vi, emacs and default keyboard bindings. You can set the binding using the keymap preference. For more information, see help cli preference. For a detailed description of the default mapping, see the Traffic Management Shell (tmsh) Reference Guide.

Note that all mappings provide command-line editing and the capability to search the command history.

WILDCARD OBJECT IDENTIFIERS

You can specify configuration object identifiers using glob and regular expression syntax.

For glob and regular expression syntax rules, see help glob and help regex. Note that you can escape the glob and regular expression special characters using a back slash.

The following examples show how to use glob and regular expressions in tmsh.

Uses a glob expression to display the configuration of all nodes that begin with 10.1..

```
list ltm node 10.1.*
```

Uses a regular expression to display the configuration of all nodes that begin with 10. and contain .44.. Note that a regular expression must begin with an @ symbol. This identifies to tmsh that the identifier should be treated as a regular expression and not a glob or standard object identifier. The leading @ is not part of the regular expression.

```
list ltm node @^10\..*\44\.
```

PREFERENCES

You can customize the behavior of tmsh. For more information, see help cli preference.

FILES

tmsh manages several files in a user's home directory.

\$HOME/.tmsh-history- contains command history.

STATISTICS

You can use tmsh to display statistics, including historical performance statistics. You can select the format in which the statistics display, as well as reset the statistics for some of the tmsh components. To determine if statistics are available for a component, see the man page for the specific component.

The following examples show how to display and reset statistics for the net interface component from the root module.

```
show net interface
reset-stats net interface
```

The following examples show how to display and reset statistics for the net interface component from the net module.

```
show interface
reset-stats interface
```

AUTOMATING TMSH

You can use tmsh to build TCL scripts to automate management of the BIG-IP. See the cli script help page.

COMMAND LINE OPTIONS

The following options can be specified when tmsh is started from the system shell.

-a tmsh does not write commands to the command history file.

Note that if auditing is enabled, tmsh continues to write commands to the audit log. This option is useful when writing scripts from the system shell, because it stops the scripts from filling up the command history file. This option applies to the non-interactive mode only.

-c Run the specified command. A command that contains multiple arguments must be in quotes. No other options

may be specified after -c

-d [ip address | host name]
Connects to the specified blade in a clustered system.

-e Disables video highlighting in tmsh.

-h Displays options you can use when accessing tmsh from the system shell.

-m Generates a tmsh debug log named tmsh.out in the current directory.

Note that when you run a tmsh script, the shell generates a debug log file for the script named tmsh.out.[script name].

Using this option causes tmsh to run significantly slower.

-q Prevents tmsh from responding to user actions with questions. This option is useful when writing non-interactive shell scripts from the system shell.

-r
This option allows the user to run TMSH the specified version. This is used to provide backwards compatibility for older TMSH syntax only. The version must be specified in the format maj.min.pt, for example 11.5.0

SEE ALSO

Detailed information on the following topics is available through the help command: cli preference, cli script, glob, help, regex, and sys provision.

For complete information about tmsh, see the Traffic Management Shell (tmsh) Reference Guide. This guide is available on the AskF5(sm) Knowledge Base ().

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2014-02-18 tmsh(1)

Commands

cd

NAME
cd command - Change the current working folder.

MODULE
All tmsh modules.

SYNTAX
Use the command cd to change the current working folder.

cd [folder name]
cd /[folder name]

DESCRIPTION
The command cd [folder name] changes the current working folder to allow the user navigation around the folder system (see sys folder). The command pwd displays the current working directory.

The current working folder may be listed in the tmos command prompt while in tmsh interactive mode (see cli preference).

Folder names are separated by a forward slash /.

There are two built-in folders:

/ is the root folder

/Common is the default folder for creating new configurations objects.

Additionally, the following directory entries:

. is the current folder

.. is the parent folder

EXAMPLES

```
cd /Common
```

Change the current working folder to /Common.

```
cd resources
```

Change the current working folder to resources. In this example the resources folder is relative to the current working folder. As an example, if the current working folder was /Common, the new working folder will be /Common/resources.

```
cd resources/profiles/udp
```

Multiple folders may be specified. Tab complete assists filling the command line with folder names.

```
cd /
```

Make the current working folder the root folder.

```
cd ../Alpha
```

Change the working directory by first going to the parent, and then switch to the sub-folder Alpha.

SEE ALSO

help, pwd, sys folder, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-08-31 cd(1)

cp

NAME

cp command - Creates a copy of a TMOS(tm) configuration object.

MODULE

All tmsh modules.

SYNTAX

Use the command cp within a tmsh module to create a copy of the component that resides in that module. To create a copy component that resides in another module, use the full path to the component.

```
cp [component] [source] [destination]
cp / [module...module] [component] [source] [destination]
```

DESCRIPTION

You must provide a unique name for each component destination of the copy operation.

EXAMPLES

```
cp template mytemplate newtemplate
```

From within the sys application module, creates a new Application Template named newtemplate with the same properties as mytemplate .

```
cp / cli script my_script1 my_script2
```

From within the sys application module, copies the my_script1 script to my_script2 within the cli module.

OPTIONS

component

Specifies the type of the component that you want to copy.

module

Specifies the module within which the component that you want to copy resides.

source

Specifies the component to be copied.

destination

Specifies a unique name for the component that will be created as part of the copy.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2010-12-06 cp(1)

create

NAME

create command - Creates a TMOS(tm) configuration component.

MODULE

All tmsh modules.

SYNTAX

Use the command create within a tmsh module to create a component that resides in that module. To create a component that resides in another module, use the full path to the component.

```
create [component] [name] [property [value]...]
create / [module...module] [component] [name] [property [value]...]
```

DESCRIPTION

You must provide a unique name for each component that you create.

EXAMPLES

```
create pool pool1
```

From within the gtm module, creates a Global Traffic Manager pool named pool1.

```
create / ltm pool my_pool
```

From within the gtm module, creates a Local Traffic Manager pool named my_pool.

OPTIONS

component

Specifies the type of the component that you want to create.

module

Specifies the module within which the component that you want to create resides.

name Specifies a unique name for the component.

property [value]...

Specifies properties for the component and their values.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 create(1)

delete

NAME

delete command - Deletes a tmsh component.

MODULE

All tmsh modules.

SYNTAX

Use the command delete within a tmsh module to delete a component that resides in that module. To delete a component that resides in another module, use the full path to the component.

```
delete [component] [name]
```

delete / [module...module] [component] [name]

DESCRIPTION

You must provide the name of the component that you want to delete.

EXAMPLES

```
delete pool pool1
```

From within the gtm module, deletes the Global Traffic Manager pool named pool1.

```
delete / ltm pool my_pool
```

From within the gtm module, deletes the Local Traffic Manager pool named my_pool.

OPTIONS

component

Specifies the type of the component that you want to delete.

module

Specifies the module within which the component that you want to delete resides.

name Specifies the name of the component that you want to delete. All may be used as an identifier for most component types.

recursive

Deletes all items in the current folder and all sub-folders that match the module, component and the name specified. all may be used as the name identifier with this command.

Note: When using recursive and all together, you will be prompted to verify this action. If you wish to disable this prompt, you may run tmsh using the -q command-line option. This is very useful when writing scripts that use this command.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 delete(1)

edit

NAME

edit command - Opens the specified components in an editor.

MODULES

All tmsh modules.

SYNTAX

Use the command edit to create components or modify the configuration of components using a text editor. To edit a component that resides in another module, use the full path to the component.

```
edit [component] [name ... name | all]
```

```
edit / [module...module] [component] [name ... name | all]
```

DESCRIPTION

You can use the command edit to create or modify components in the auth, cli, gtm, ltm, net, sys and wom modules, and iRules(r).

If you are assigned the role of Administrator, when you use the command edit, the system starts the vi editor. If you are assigned any other role, the system starts the pico/nano editor.

The system saves, in a temporary directory, the text file, named data, that you are editing. When you save the file and close the editor, the system checks for errors, and then prompts you with an opportunity to continue editing and resolve any errors.

When you edit an existing component that can have associations, such as a Global Traffic Manager wide IP that can have pool member associations. but the component does not currently have associations, to create the new associations, you must use the full command syntax in the text file. For the full command syntax for each component, see the associated man page.

When you edit a component that has associations with components that are children of the component you are editing, the text file contains a line for the configuration of the child components that begins with the command modify, for example: pools modify { [existing pool members configurations] }. In this case, if you want to add or delete pool members, you must add additional lines to the text file, for example: pools delete { [pool members to delete] }.

If you want the text file that opens to contain all of the editable properties of the component that you want to edit, you must use the all-properties option at the end of the edit command sequence; otherwise, only the non-default properties display in the text file.

EXAMPLES

```
edit / gtm pool a*
```

From the root module, opens a file in an editor in which you can modify the configuration of all Global Traffic Manager pools with names that start with the letter a using the template that displays in the editor.

```
edit datacenter new_dc
```

From the gtm module, opens a file in an editor in which you can create the Data Center named new_dc using the template that displays in the editor.

```
edit datacenter a*
```

From the gtm module, opens a file in an editor in which you can edit all existing datacenters with names that begin with the letter a.

```
edit datacenter new_datacenter existing_datacenter
```

From the gtm module, opens a file in an editor in which you can create a new datacenter and edit an existing datacenter. Note that when the file opens, a template displays that you can use to create a new datacenter followed by the configuration of the existing datacenter.

```
edit rule rule_1
```

From the gtm module, opens a file in an editor in which you can create an iRule named rule_1 using the template that displays in the editor.

When the editor opens, and you are creating or editing an iRule, you must enclose the iRule syntax in brackets, for example, [...iRule...]. Note that the template includes the brackets.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

component

Specifies the type of component that you want to create or modify.

module

Specifies the module within which the component resides.

name Specifies a unique name of each component that you want to create or modify.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-03-22 edit(1)

exit

NAME

exit command - Exits a tmsh module or component.

MODULE

All tmsh modules.

SYNTAX

Use the command exit within a tmsh module or component to leave that module or component and return to the higher level of the shell structure.

```
exit
```

Note that to exit tmsh and return to the BIG-IP(r) system prompt, use the command quit.

DESCRIPTION

For more information about the structure of tmsh, see the Traffic Management Shell (tmsh) Reference Guide.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-04-05 exit(1)

generate

NAME

generate - Generate signed scripts using different algorithms for components (for example, iRules).

MODULE

All tmsh modules.

DESCRIPTION

Use the generate command to generate signed scripts for components. Currently two algorithms are supported: checksum and signature.

```
generate checksum
generate signature signing-key
```

SEE ALSO

ltm rule, sys application template

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2014-04-08 generate(1)

help

NAME

help command - Displays context-sensitive help text.

MODULE

All tmsh modules.

SYNTAX

Use the command help within a tmsh module to display information about the components that reside within that module, or at the component level to display help about the component. To display help for a component that resides in one module from within another module, use the full path to the component.

Type the question mark (?) character anywhere in tmsh to display a list of modules, components, and commands that are available within the module in which you are currently working.

```
?
help
help [module...module]
help [component]
help / [module...module] [component]
help search [text]
```

DESCRIPTION

You can display tmsh man pages using the command help.

EXAMPLES

?

From within the gtm module, displays a list of modules, components, and commands that are available.

```
help pool
```

From within the gtm module, displays help about Global Traffic Manager pools.

```
help / ltm pool
```

From within the gtm module, displays help about Local Traffic Manager pools.

OPTIONS

component

Specifies the type of the component for which you want to display help.

search

Use the search option to find help topics that contain the specified text. The search is case insensitive. Text that contains a space or special tmsh characters must be quoted. Note that the search will not always find text that spans multiple lines.

module

Specifies the module within which the component for which you want to display help resides.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2012-10-19 help(1)

install

NAME

install - Install and update components.

MODULE

All tmsh modules.

DESCRIPTION

Use the command install to install or update the following components. For the description and syntax see the help page for each component.

sys license

sys software block-device-hotfix

sys software block-device-image

sys software hotfix

sys software image

SEE ALSO

sys license, sys software block-device-hotfix, sys software block-device-image, sys software hotfix, sys software image, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2014-04-20 install(1)

list

NAME

list command - Displays components that you have permission to view.

MODULE

All tmsh modules.

SYNTAX

Use the list command within a tmsh module to display the properties of the components in that module. To display the properties of the components in one module from within another module, use the full path to the component.

list [component]

list [component] [name]
list [component] [name] [property]
list / [module...module] [component] [name] [property]
options:
 all-properties
 current-module
 non-default-properties
 one-line
 partition
 recursive

DESCRIPTION

When the default Read partition is All, use the list command to display all of the components that you have permission to view within a tmsh module. When you specify a Read partition, the list command displays:

- Â· Only the components that you have permission to view in the current partition
- Â· All of the components that are not in partitions
- Â· All of the components in partition Common

EXAMPLES

list / ltm

From within the gtm module, displays the properties of all of the components in the ltm module, including the components in the ltm monitor, ltm persistence, and ltm profile modules.

list / ltm current-module

From within the gtm module, displays the properties of all of the components in the ltm module, not including the components in the ltm monitor, ltm persistence, and ltm profile modules.

list pool

From within the gtm module, displays the properties of all of the Global Traffic Manager pools.

list pool all-properties

From within the gtm module, displays all of the properties of all of the Global Traffic Manager pools.

list pool monitor

From within the gtm module, displays the monitor associated with each Global Traffic Manager pool.

list / ltm pool

From within the gtm module, displays the properties of all of the Local Traffic Manager pools.

OPTIONS

all-properties

Displays the values of all of the properties of the specified component.

component

Specifies the component that you want to display.

current-module

Specifies to display only the components that reside in the specified module, not the components that reside in the sub-modules of that module.

For example, from within the ltm module to display only the components in the gtm module, and not the components in the gtm monitor and gtm settings sub-modules, use the following command sequence: list / gtm current-module.

module

Specifies the module within which the component that you want to display resides.

Note: When you use the command list at the module level, by default, the system does not display all of the components that reside in the specified module. To display the properties of some components you must explicitly specify the component. For example, from the ltm module, to display the virtual addresses for the Local Traffic Manager, use this command sequence:

list virtual-address

For more information about displaying the properties of a component, see the man page for the component.

name Specifies the unique name of the component.

non-default-properties

Displays the values of all of the properties for which a user changed the value from the default value for the specified component.

one-line

Displays the configuration for each object on one line. Configuration that consists of scripts will not be formatted on to a single line. This include ltm and gtm iRules and tmsh scripts.

partition

Displays the administrative partition within which the specified component exists.

property

Specifies the property of the component that you want to display.

recursive

Specifies to display the components not only from the current folder but also from all sub-folders recursively.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 list(1)

load

NAME

load command - Replaces the running configuration of the BIG-IP(r) system with the configuration in the specified files. You can also use this command to import an ASM policy from a file / standard input, and to install the Anti-fraud engine / signatures update.

MODULE

All tmsh modules.

SEE ALSO

save, tmsh, asm policy, ltm dns dns-express db, sys config, sys geoip, sys ucs

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2014-12-30 load(1)

modify

NAME

modify command - Modifies a tmsh component.

MODULE

All tmsh modules.

SYNTAX

Use the command modify within a tmsh module to modify a component that resides in that module. To modify a component in one module from within another module, use the full path to the component.

modify [component] [name] [property [value]]...

modify / [module...module] [component] [name] [property [value]]...

DESCRIPTION

You must provide the name of the component that you want to modify.

You can apply one or more property settings to multiple components using a single command sequence. For example, to associate the Local Traffic Manager pool named pool-1 with the virtual servers named virtual-1 and virtual-2, use this command sequence: modify ltm virtual virtual-1 virtual-2 pool pool-1

EXAMPLES

modify pool pool1 disabled

From within the gtm module, disables the Global Traffic Manager pool named pool1.

modify / ltm pool my_pool disabled

From within the gtm module, disables the Local Traffic Manager pool named my_pool.

OPTIONS

component

Specifies the type of the component that you want to modify.

module

Specifies the module within which the component that you want to modify resides.

name Specifies the unique name of the component that you want to modify.

property [value]...

Specifies the properties of the component that you want to modify and their new values.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 modify(1)

mv

NAME

mv command - Renames or moves a TMOS(tm) configuration object.

MODULE

All tmsh modules.

SYNTAX

Use the mv command within a tmsh module to move or rename the component that resides in that module. To move a component that resides in another module, use the full path to the component.

```
mv [component] [source] [destination]
```

```
mv / [module...module] [component] [source] [destination]
```

DESCRIPTION

You must provide a unique name for the source and destination of the move operation.

WARNING Currently MV is an experimental feature. By using this feature, you may be subject to loss of statistics and disruption in GTM service. If you plan to move or rename a Virtual Server, please contact your GTM administrator before doing so. You may enable this feature by setting the appropriate db variable. This can be done by issuing the command:

```
modify /sys db mcpd.mvenabled value true
```

This will turn on the feature and allow moving and rename of select objects through TMSH only. Once you have finished using the feature, we recommend disabling it once again. You may do this by issuing the following command:

```
modify /sys db mcpd.mvenabled value false
```

Please use responsibly.

EXAMPLES

```
mv cm device bigip seattle32
```

Renames the device named bigip to seattle32.

```
mv ltm pool mypool myotherpool
```

Renames the LTM Pool named mypool to myotherpool.

```
mv ltm pool /Common/by/mypool /Common/myotherpool /Common/sub/mythirdpool to-folder /Partition2/sub1
```

Moves the 3 pools in 3 different locations named mypool, myotherpool and mythirdpool into a single folder in another partition.

OPTIONS

to-folder

Specifies the folder to move the item or items into.

component

Specifies the type of the component that you want to move.

destination

Specifies a unique name for the component.

module

Specifies the module within which the component that you want to move resides.

source

Specifies the component to be moved.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012. All rights reserved.

BIG-IP 2014-03-25 mv(1)

publish

NAME

publish - Finalizes changes in the policy by creating a read-only copy of it.

MODULE

All tmsh modules.

DESCRIPTION

Use the command publish to make wam policies available for usage in wam applications. You can also use this command to apply asm policies. For the description and syntax see the help page for wam policy or asm policy.

SEE ALSO

asm policy, wam policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-09-05 publish(1)

pwd

NAME

pwd command - Display the current working folder.

MODULE

All tmsh modules.

SYNTAX

Use the command pwd to display the current working folder.

pwd

DESCRIPTION

Display the current working folder

EXAMPLES

pwd

SEE ALSO

cd, help, sys folder, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

quit

NAME
quit command - Exits tmsh.

MODULE
All tmsh modules.

SYNTAX
Use the following command at the tmsh prompt to close tmsh and return to the BIG-IP(r) system prompt.

quit

Note that to exit a tmsh module or component, you use the command exit.

SEE ALSO
tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

reboot

NAME
reboot command - Reboots the system or boots the system into a different volume.

MODULE
All tmsh modules.

SYNTAX
reboot
options:
slot [[slot number] | all]
volume [name]

DESCRIPTION
You can use the command reboot to reboot the system or cluster. If you do not specify an option, the local system reboots.

You can use the volume option to reboot a system into a specific volume. For a cluster, you can use the volume option to reboot all slots into the specified volume.

Additionally, for a cluster, you can use the slot option to reboot either a specific slot or all slots. Note that the slot option does not modify the active volume.

EXAMPLES
reboot

Immediately reboots the running image.

reboot volume HD1.2

If the volume HD1.2 has a complete image on it, the system (or cluster) reboots into that image immediately. However, if a software installation is in progress on the volume the system reboots as soon as the installation is complete.

If the volume contains software that is not a version permitted by the license a warning will be displayed requiring the user to input Y/N with the Y standing for 'Yes', proceed with the reboot, or N for 'No', stop the reboot and return to the tmsh command line.

OPTIONS
slot [[slot number] | all]
Reboots either a specific slot or all slots in a cluster, without changing the active volume of the

slot(s).

This option is only available in a clustered environment.

Note: The slot and volume options are mutually exclusive.

volume

Specifies the volume that you want to boot. The volume you specify becomes the default boot volume. You cannot specify the active volume. In a clustered environment all slots reboot into the same volume.

Note: The slot and volume options are mutually exclusive.

SEE ALSO

install, sys software hotfix, sys software image, sys software status, sys software volume, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2019-05-02 reboot(1)

reset-stats

NAME

reset-stats - Resets statistics for the specified components.

MODULE

All tmsh modules.

SYNTAX

Use the command reset-stats within a tmsh module to reset the statistics for the specified component to zero. To reset the statistics for the specified component in one module from within another module, use the full path to the component.

```
reset-stats [component]
reset-stats [component] [name]
reset-stats / [module...module] [component]
reset-stats / [module...module] [component] [name]
```

DESCRIPTION

You can reset statistics for a group of components, or you can reset statistics for a specific component.

After you reset statistics, when you run the command show, you may see a value of nan. This stands for not a number, which indicates that no data is currently available. Wait a few moments and run the command show again, and in most cases the nan value will be replaced by an integer value.

It is important to note the following when you reset statistics:

- Â· For a data center, the system also resets the statistics for the servers in that data center.
- Â· For a Global Traffic Manager server, the system also resets the statistics for the virtual servers on that server.
- Â· For a Global Traffic Manager pool, the system also resets the statistics for the pool members.
- Â· For a Local Traffic Manager pool, the system also resets the statistics for the pool members.
- Â· For a VLAN, you must reset the statistics for the trunks and interfaces associated with the VLAN.
- Â· You cannot reset statistics for system-supplied profiles.

EXAMPLES

```
reset-stats pool
```

From within the gtm module, resets the statistics for all of the Global Traffic Manager pools.

```
reset-stats pool pool1
```

From within the ltm module, resets the statistics for the Local Traffic Manager pool named pool1.

```
reset-stats / ltm pool my_pool
```

From within the gtm module, resets the statistics for the Local Traffic Manager pool named my_pool.

```
reset-stats all-stats
```

From within the sys performance module, resets all performance statistics for the system.

OPTIONS

component

Specifies the type of the component for which you want to reset statistics.

module

Specifies the module within which the component for which you want to reset statistics resides.

name Specifies the unique name of the component for which you want to reset statistics.

SEE ALSO

tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-04-05 reset-stats(1)

restart

NAME

restart command - Restarts a service on the BIG-IP(r) system.

MODULE

All tmsb modules.

SYNTAX

Use the command restart within tmsb to restart a specified service.

restart

options:

/sys service [service name]

DESCRIPTION

You can use the command restart to restart a specified service.

EXAMPLES

restart /sys service mcpd

Restarts the mcpd daemon.

restart /sys service snmpd

Restarts the snmpd daemon.

OPTIONS

Tip: Use the tab completion feature to see a list of available services.

SEE ALSO

start, stop, sys service, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 restart(1)

run

NAME

run command - Runs the specified program.

MODULE

All tmsb modules.

SYNTAX

Use the run command within tmsh to run a specified utility.

run

options:

```
/cli script [arguments]
/cm add-to-trust
/cm config-sync
/cm remove-from-trust
/cm sniff-updates
/cm watch-devicegroup-device
/cm watch-sys-device
/cm watch-trafficgroup-device
/gtm big3d_install [arguments]
/gtm bigip_add [arguments]
/gtm gtm_add [arguments]
/ltn monitor [arguments]
/security anti-fraud engine-update
/security anti-fraud signatures-update [arguments]
/sys air-filter-reset
/util bash [arguments]
/util diadb [arguments]
/util dig [arguments]
/util dnat [arguments]
/util get-ccn-dossier
/util get-dossier [arguments]
/util ihealth [arguments]
/util qkcloud [arguments]
/util ipsecalgdb [arguments]
/util lsndb [arguments]
/util netstat [arguments]
/util ping [arguments]
/util ping6 [arguments]
/util qkview [arguments]
/util racoonctl [arguments]
/util sipdb [arguments]
/util ssh-keyswap [arguments]
/util sys-icheck [arguments]
/util tcpdump [arguments]
/util tracepath [arguments]
/util tracepath6 [arguments]
/util traceroute [arguments]
/util traceroute6 [arguments]
/wom diagnose-conn
/wom verify-config
```

DESCRIPTION

You can use the run command to run the specified program, utility or process.

You can read about the arguments that are available for the utilities in the cm module using the following command sequence:

```
help /cm [utility name]
```

You can read about the arguments that are available for the utilities in the gtm module using the following command sequence:

```
help /gtm [big3d_install | bigip_add | gtm_add]
```

You can read about the arguments that are available for the utilities in the ltm module using the following command sequence:

```
help /ltn monitor [type]
```

You can read about the arguments that are available for the utilities in the util module using the following command sequence:

```
help /util [utility name]
```

Note: Some tmsh features, such as tab completion, context-sensitive help, paging, and grep, are not available for utilities.

When you are building a batch mode transaction in tmsh, if you type the run command, the system runs the specified program immediately. It does not add the run command to the transaction that you are building.

EXAMPLES

```
help /util ping
```

Displays the help page for the ping utility.

OPTIONS

```
/cli script
```

Run the specified script, with provided arguments.

```
/cm add-to-trust
```

Add a device to a trust domain.

`/cm config-sync`

Synchronize the configuration between devices.

`/cm remove-from-trust`

Remove a device from a trust domain.

`/cm sniff-updates`

Display the commit ID updates that occur over the CMI communications channel. When you troubleshoot a ConfigSync issue, it is helpful to determine which device group member has the latest commit ID update and contains the most recent configuration. You can then decide whether to replicate the newer configuration to the group, or perform a ConfigSync operation that replicates an older configuration to the group, thus overwriting a newer configuration.

`/cm watch-devicegroup-device`

Display information about the devices in the device group to which the local device belongs.

`/cm watch-sys-device`

Display information about the local device.

`/cm watch-trafficgroup-device`

Display information about the traffic groups associated with devices in a device group.

`/gtm big3d_install`

Specifies to install the big3d daemon.

`/gtm bigip_add`

Specifies the BIG-IP systems that you want to add to the Global Traffic Manager configuration.

`/gtm gtm_add`

Specifies the Global Traffic Manager systems that you want to add to the Global Traffic Manager configuration.

`/ltm monitor`

Performs a one-shot test of a custom LTM health monitor against a specified target node.

`/security anti-fraud engine-update`

For the description and syntax see the help page for security anti-fraud engine-update.

`/security anti-fraud signatures-update`

For the description and syntax see the help page for security anti-fraud signatures-update.

`/sys air-filter-reset`

Runs the command to reset the timer and the retry counter for air filter notifications.

`/util bash`

Accesses the system shell.

`/util diadb`

Displays Diameter persistence entries. The diadb utility displays diameter persistence entries or delete a particular persistence entry.

`/util dig`

Runs the specified dig command. The dig utility queries DNS name servers.

`/util dnat`

Runs the specified dnat command for the purpose of doing forward/reverse mapping of addresses for DNAT.

`/util get-ccn-dossier`

Runs the get_ccn_dossier utility for the purpose of displaying system information for dossier creation.

`/util get-dossier`

Runs the get_dossier utility for the purpose of displaying system license dossier information.

`/util ihealth`

Runs the ihealth utility for the purpose of uploading a new or existent qkview file to ihealth.

`/util qkcloud`

Runs the qkcloud utility for the purpose of displaying cloud meta-data.

`/util ipsecalgdb`

Runs the ipsecalgdb utility to view IPsecALG translation, and pending IKE connection count entries.

`/util lsndb`

Runs the lsndb utility to view Large Scale NAT persistence entries, inbound mappings, client connection counts, and PCP mappings.

`/util netstat`

Displays network connections, routing tables, interface statistics, masquerade connections, and multicast memberships.

`/util ping`

Runs the specified ping command. The ping utility sends ICMP echo requests to network hosts.

`/util ping6`

Runs the specified ping6 command. The ping6 utility sends ICMPv6 echo requests to network hosts.

`/util qkview`

Runs the specified qkview command. The qkview utility gathers diagnostic information from a BIG-IP system.

`/util racoonctl`

Runs the specified racoonctl command. The racoonctl utility is used to control operation of the racoon daemon.

`/util sipdb`

Displays SIP persistence entries. The sipdb utility displays specific persistence entries and deletes a particular persistence record.

`/util ssh-keyswap`

Runs the keyswap.sh script for managing SSH keys on the BIG-IP.

`/util sys-icheck`

Runs the specified sys-icheck command. The sys-icheck utility verifies all RPM packages and files.

`/util tcpdump`

Runs the specified tcpdump command. The tcpdump utility prints headers and content of network traffic.

`/util tracepath`

Displays the route packets take to a network host.

`/util tracepath6`

Displays the route packets take to an IPv6 network host.

`/util traceroute`

Displays the route packets take to a network host.

`/util traceroute6`

Displays the route packets take to an IPv6 network host.

`/wom diagnose-conn`

Runs the specified diagnose-conn script, which detects the sources of network connection and performance problems in a WAN optimization configuration.

`/wom verify-config`

Runs the specified verify-config script, which detects errors in the configuration of the WAN Optimization Manager.

SEE ALSO

cli script, cm add-to-trust, cm config-sync, cm remove-from-trust, cm sniff-updates, cm watch-devicegroup-device, cm watch-sys-device, cm watch-trafficgroup-device, gtm big3d_install, gtm bigip_add, gtm gtm_add, ltm monitor, security anti-fraud engine-update, security anti-fraud signatures-update, util bash, util diadb, util dig, util dnat, util ihealth, util qkcloud, util ipsecalgdb, util lsndb, util netstat, util ping, util ping6, util qkview, util racoonctl, util sipdb, util ssh-keyswap, util sys-icheck, util tcpdump, util tracepath, util tracepath6, util traceroute, util traceroute6, wom diagnose-conn, wom verify-config

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015, 2017. All rights reserved.

BIG-IP 2018-04-16 run(1)

save

NAME

save - Writes the running configuration of the BIG-IP(r) system to the specified file.

MODULE

All tmsh modules.

DESCRIPTION

You can use the save command to write changes that you make to the running configuration of the BIG-IP system to the specified file. You can also use this command to save an analytics report to a file on the BIG-IP(r) system or to export an ASM policy to a file / standard output.

SEE ALSO

analytics report, asm policy, load, sys config, sys ucs, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

send-mail

NAME

send-mail - Send an e-mail to a list of recipients containing configuration or statistical information about the BIG-IP(r) system.

MODULE

All tmsh modules.

DESCRIPTION

You can use the send-mail command to send an analytics report from the BIG-IP system to a list of e-mail recipients.

SEE ALSO

analytics report, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 send-mail(1)

show

NAME

show command - Displays statistics for and the status of specified components.

MODULE

All tmsh modules.

SYNTAX

Use the show command within a tmsh module to display statistics for and the status of components in that module. To display statistics for and the status of components in another module, use the full path to the component.

show

show [component]

show [component] [name]

show / [module] [component] [name]

options:

all-stats

current-module

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

(detail | global | historical)

field-fmt

running-config

recursive

DESCRIPTION

You can use the show command to specify the unit value in which the system displays statistics and the type of statistics that you want the system to display.

After you reset statistics, when you run the command show, you may see a value of nan. This stands for not a number, which indicates that no data is currently available. Wait a few moments and run the show command again, and in most cases the nan value is replaced by an integer value. For more information, see help reset-stats.

EXAMPLES

show / ltm current-module

From within the gtm module, displays statistics and status for all the components within the ltm module, but not the components in the ltm monitor, ltm persistence, and ltm profile modules.

show pool

From within the gtm module, displays statistics and status for all Global Traffic Manager pools.

show pool pool1

From within the gtm module, displays statistics and status for the Global Traffic Manager pool named pool1.

show / ltm pool

From within the gtm module, displays statistics and status for all Local Traffic Manager pools.

show / ltm profile tcp global

From within the gtm module, displays global statistics and status for all Local Traffic Manager TCP profiles in the system default unit.

OPTIONS

all-stats

Displays all of the available system performance statistics.

component

Specifies the type of the component for which you want to show statistics and status.

current-module

Specifies to display only the components that reside in the specified module, not the components that reside in the sub-modules of that module.

For example, from within the ltm module to display only the components in the gtm module, and not the components in the gtm monitor and gtm settings sub-modules, use this command sequence: show / gtm current-module.

default

Displays data in the simplest units. For example, if the value of the data is 1,200,001, the system displays 1.20M; however, if the value of the data is 1,200, the system displays 1.2K.

detail

Displays detailed data for the specified component and associated components. Note that this option is available for only a partial set of tmsh components.

You can use the tab completion and context-sensitive help features to determine if this option is available. For more information about these features, see help.

field-fmt

Displays data as a list of options and their values. The option names can be used to retrieve statistics and status values in a shell script, see cli script.

gig Displays data in parts per billion.

global

Displays global statistics for the specified component that includes statistics for all components of the specified type. Note that this option is available for only a partial set of tmsh components. You can use the tab completion and context-sensitive help features to determine if this option is available.

historical

Displays historical statistics for the specified component. Note that this option is available only for a partial set of tmsh components. You can use the tab completion and context-sensitive help features to determine if this option is available.

kil Displays data in parts per thousand.

lines

Specifies how many lines of the log that you want the system to display.

meg Displays data in parts per million.

module

Specifies the module within which the component for which you want to show statistics and status resides.

Note: When you use the command show at the module level, by default, the system does not display all of the components that reside in the module. To display some components you must explicitly specify the component. For example, from the ltm module, to display the statistics for and status of the virtual addresses of the Local Traffic Manager, use the following command sequence:

show virtual-address

For more information about displaying statistics for and status of a component, see the man page for the component.

name Specifies the unique name of the component for which you want to show statistics and status.

range

Specifies a date range for the logs that you want the system to display, for example:

2d-4d

Specifies 2 - 4 days ago.

3d Specifies 3 days ago to now.

epoch--7/25:12:00:00

Specifies everything older than July 25th at noon.

2008-07-25--2008-07-28:13:30
Specifies between July 25th and 28th at 1:30 p.m.

raw Displays raw data.

recursive
Specifies to display the components not only from the current folder, but also from all sub-folders recursively.

running-config
Displays the running configuration of the components that you have permission to view within a tmsh module, if the default Read partition is All. If you specify a Read partition, this option displays only the components that you have permission to view in the current partition, all of the components that are not in partitions, and all of the components in partition Common. Note that this option is valid only for tmsh components you can configure.

The running-config option must be specified immediately after the show command, for example:

```
show running-config ltm pool
```

SEE ALSO
cli script, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 show(1)

shutdown

NAME
shutdown command - Shuts down the system.

MODULE
All tmsh modules.

SYNTAX
shutdown
options:
slot [[slot number] | all]

DESCRIPTION
You can use the command shutdown to power down the system or cluster. If you do not specify an option, the local system shuts down.

For a cluster, you can use the slot option to shut down either a specific slot or all slots.

EXAMPLES
shutdown

Immediately shuts down the running system.

OPTIONS
slot [[slot number] | all]
Shuts down either a specific slot or all slots in the cluster. This option is only available in a clustered environment.

SEE ALSO
reboot, install

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-03-21 shutdown(1)

start

NAME

start command - Starts a service on the BIG-IP(r) system.

MODULE

All tmsh modules.

SYNTAX

Use the start command within tmsh to restart a specified service.

start

options:

/sys service [service name]

DESCRIPTION

You can use the start command to start a specified service.

EXAMPLES

start /sys service mcpd

Starts the mcpd daemon.

start /sys service snmpd

Starts the snmpd daemon.

OPTIONS

Tip: Use the tab completion feature to see a list of available services.

SEE ALSO

restart, stop, sys service, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 start(1)

stop

NAME

stop command - Stops a service or test operation that is running on the BIG-IP(r) system.

MODULE

All tmsh modules.

SYNTAX

Use the command stop within tmsh to stop a running service or test operation.

stop

options:

/sys service [service name]

/ltm monitor [arguments]

DESCRIPTION

You can use the command stop to stop a running service or test operation.

EXAMPLES

stop /sys service mcpd

Stops the mcpd daemon.

stop /sys service snmpd

Stops the snmpd daemon.

stop /ltm monitor http my_http

Cancels a pending health monitor test for the custom ltm http monitor my_http.

OPTIONS

Tip: Use the tab completion feature to see a list of available services.

SEE ALSO

ltm monitor, restart, start, sys service, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013, 2016-2017. All rights reserved.

BIG-IP 2017-03-24 stop(1)

submit

NAME

submit - Runs the transaction that you are creating.

MODULE

All tmsh modules.

SYNTAX

Use the submit command to run a transaction that you are creating.

submit transaction

DESCRIPTION

You can use the submit command to run a transaction, which is a series of commands that you enter in transaction mode.

For more information about creating transactions, see cli transaction.

SEE ALSO

cli transaction, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-04-05 submit(1)

Modules

analytics

analytics afm-sweeper report

NAME

report - Displays an afm-sweeper analytics report.

MODULE

analytics afm-sweeper

SYNTAX

Show, save or send an analytics afm-sweeper report using the syntax shown in the following sections.

DISPLAY

show report view-by [policy-name | context-type | context-name | eviction-reason | client-ip]

options:

drilldown {

{

entity [policy-name | context-type | context-name | eviction-reason | client-ip]

values

{

[value ...]

}

} ...

}

field-fmt

```

include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SAVE

save report view-by [policy-name | context-type | context-name | eviction-reason | client-ip]

options:

```

drilldown {
  {
    entity [ policy-name | context-type | context-name | eviction-reason | client-ip ]
    values
    {
      [value ...]
    }
  } ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SEND

send-mail report view-by [policy-name | context-type | context-name | eviction-reason | client-ip]

options:

```

drilldown {
  {
    entity [ policy-name | context-type | context-name | eviction-reason | client-ip ]
    values
    {
      [value ...]
    }
  } ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate afm-sweeper analytics reports. You can generate an afm-sweeper analytics report for the following entities:

Â· policy-name - Name of the sweeper policy.

Â· context-type - The context type of the eviction. Can be one of three possible values: "Virtual Server", "Route Domain", "Global".

Â· context-name - The name of the context. Depends on the context-type this value can be either a virtual server name (e.g. "my_vs1"), route domain id (e.g. 317) or "Global".

Â· eviction-reason - The algorithm used by sweeper that caused the eviction of a connection.

Â· client-ip - The source IP of the evicted connection.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics afm-sweeper report view-by client-ip
```

```
show analytics afm-sweeper report view-by client-ip drilldown { { entity context-type values { Global } } }
```

```
send-mail analytics afm-sweeper report view-by client-ip measures { avg-connections-age } limit 15 order-by { { measure avg-connections-age sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-connections-age

The average time that the connections lasted before evicted for the selected filter (entity).

evicted-connections

The total number of evicted connections for the selected filter (entity).

total-bytes-in

The total number of bytes received for the selected filter (entity).

total-bytes-out

The total number of bytes sent for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

analytics afm-sweeper scheduled-report

NAME

scheduled-report - Configure scheduled reports for AFM sweeper.

MODULE

analytics afm-sweeper

SYNTAX

Configure the scheduled-report component within the analytics afm-sweeper module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the AFM sweeper module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the AFM sweeper scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics afm-sweeper report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics afm-sweeper report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics afm-sweeper report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics afm-sweeper scheduled-report(1)

analytics application-security-anomalies report

NAME

report - Displays an application-security-anomalies analytics report.

MODULE

analytics application-security-anomalies

SYNTAX

Show, save or send an analytics application-security-anomalies report using the syntax shown in the following sections.

DISPLAY

show report view-by [anomaly-type | application | policy | virtual]

options:

drilldown {

{

entity [anomaly-type | application | policy | virtual]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

```

measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SAVE

```

save report view-by [ anomaly-type | application | policy | virtual ]
options:
  drilldown {
    {
entity [ anomaly-type | application | policy | virtual ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SEND

```

send-mail report view-by [ anomaly-type | application | policy | virtual ]
options:
  drilldown {
    {
entity [ anomaly-type | application | policy | virtual ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate application-security-anomalies analytics reports. You can generate an application-security-network analytics report for the following entities:

- Â· anomaly-type - Anomaly type (Brute Force/Web Scraping)
- Â· application - Application services.
- Â· policy - Security policy.
- Â· virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics application-security-network report view-by application
```

```
show analytics application-security-network report view-by application drilldown { { entity virtual values { my_vip } } }
```

```
send-mail analytics application-security-anomalies report view-by virtual measures { rejected-requests } limit 20 order-by { { measure rejected-requests sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

rejected-requests

The total number of rejected requests for the selected filter (entity).

total-attacks

The total number of attacks for the selected filter (entity).

total-violations

The total number of violations for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-10-15 analytics application-security-anomalies report(1)

NAME

scheduled-report - Configure scheduled reports for application security anomalies (ASM anomalies).

MODULE

analytics application-security-anomalies

SYNTAX

Configure the scheduled-report component within the analytics application-security-anomalies module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the application security anomalies (ASM anomalies) module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the application security anomalies (ASM anomalies) scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics application-security-anomalies report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics application-security-anomalies report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics application-security-anomalies report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 201analytics application-security-anomalies scheduled-report(1)

analytics application-security-incidents report

NAME

report - Displays an application-security-anomalies analytics report.

MODULE

analytics application-security-anomalies

SYNTAX

Show, save or send an analytics application-security-anomalies report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ anomaly-type | application | policy | virtual ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ anomaly-type | application | policy | virtual ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
field-fmt
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

range [date range]

SAVE

save report view-by [anomaly-type | application | policy | virtual]

options:

drilldown {

{
entity [anomaly-type | application | policy | virtual]

values

{
[value ...]

} ...

}
file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{
measure [measure name]
sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [anomaly-type | application | policy | virtual]

options:

drilldown {

{
entity [anomaly-type | application | policy | virtual]

values

{
[value ...]

} ...

}

email-addresses {

[email address ...]

}

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{
measure [measure name]
sort-type [asc / desc]

} ...

}

range [date range]

smtp-config-override [smtp configuration object name]

DESCRIPTION

Use this command to generate application-security-anomalies analytics reports. You can generate an application-security-network analytics report for the following entities:

Â· anomaly-type - Anomaly type (Brute Force/Web Scraping)

Â· application - Application services.

Â· policy - Security policy.

Â· virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

show analytics application-security-network report view-by application

show analytics application-security-network report view-by application drilldown { { entity virtual values { my_vip } } }

send-mail analytics application-security-anomalies report view-by virtual measures { rejected-requests } limit 20 order-by { { measure rejected-requests sort-type desc } } format pdf email-addresses { some.one@someaddress.com }

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

rejected-requests

The total number of rejected requests for the selected filter (entity).

total-attacks

The total number of attacks for the selected filter (entity).

total-violations

The total number of violations for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, Itm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-01-20 analytics application-security-incidents report(1)

analytics application-security-network report

NAME

report - Displays an application-security-network analytics report.

MODULE

analytics application-security-network

SYNTAX

Show, save or send an analytics application-security report using the syntax shown in the following sections.

DISPLAY

show report view-by [application | virtual | request-type | policy]

options:

```
drilldown {
  {
entity [ application | virtual | request-type | policy ]
values
{
  [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SAVE

save report view-by [application | virtual | request-type | policy]

options:

```
drilldown {
  {
entity [ application | virtual | request-type | policy ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SEND

send-mail report view-by [application | virtual | request-type | policy]

options:

```
drilldown {
  {
entity [ application | virtual | request-type | policy ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate application-security-network analytics reports. You can generate an application-security-network analytics report for the following entities:

- application - Application services.
- virtual - Virtual servers.
- request-type - Request types (Legal/Alarmed/Blocked).
- policy - Security policy.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics application-security-network report view-by violation
```

```
show analytics application-security-network report view-by violation drilldown { { entity severity values { Error } } }
```

```
send-mail analytics application-security-network report view-by virtual measures {events} limit 20 order-by { { measure events sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

events

The total number of events (requests) for the selected filter (entity).

throughput

The average throughput (bits/s) for the selected filter (entity).

tps The average number of transactions per second for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, Itm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-09-08 analytics application-security-network report(1)

analytics application-security-network scheduled-report

NAME

scheduled-report - Configure scheduled reports for application security network (ASM network).

MODULE

analytics application-security-network

SYNTAX

Configure the scheduled-report component within the analytics application-security-network module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the application security network (ASM network) module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the application security network (ASM network) scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics application-security-network report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics application-security-network report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics application-security-network report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-analytics application-security-network scheduled-report(1)

analytics application-security report

NAME

report - Displays an application-security analytics report.

MODULE

analytics application-security

SYNTAX

Show, save or send an analytics application-security report using the syntax shown in the following sections.

DISPLAY

show report view-by [application | virtual | request-type | severity | rating | username | attack-type | ip-address-intelligence | policy response-code | ip | violation | country | method | protocol | session-id | url | virus]

options:

drilldown {

{

entity [application | virtual | request-type | severity | rating | username | attack-type | ip-address-intelligence | policy response-code | ip | violation | country | method | protocol | session-id | url | virus]

values

{

```

[value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SAVE

```

save report view-by [ application | virtual | request-type | severity | rating | username | attack-type | ip-address-intelligence | policy
  response-code | ip | violation | country | method | protocol | session-id | url | virus ]

```

options:

```

drilldown {
  {
entity [ application | virtual | request-type | severity | rating | username | attack-type | ip-address-intelligence | policy
response-code | ip | violation | country | method | protocol | session-id | url | virus ]

```

values

```

{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SEND

```

send-mail report view-by [ application | virtual | request-type | severity | rating | username | attack-type | ip-address-intelligence | policy
  response-code | ip | violation | country | method | protocol | session-id | url | virus ]

```

options:

```

drilldown {
  {
entity [ application | virtual | request-type | severity | rating | username | attack-type | ip-address-intelligence | policy
response-code | ip | violation | country | method | protocol | session-id | url | virus ]

```

values

```

{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate application-security analytics reports. You can generate an application-security analytics report for the following entities:

- application - Application services.

- virtual - Virtual servers.
- request-type - Request types (Legal/Alarmed/Blocked).
- severity - Violation severities.
- rating - Violation ratings.
- username - User names.
- attack-type - Attack type of the illegal request.
- ip-address-intelligence - IP Address reputation categories.
- policy - Security policy.
- response-code - Response codes.
- ip - Source IP addresses.
- violation - Violation types.
- country - Countries of the source IP address.
- method - HTTP methods.
- protocol - Protocols (HTTP/HTTPS).
- session-id - IDs of sessions.
- url - Requested URLs.
- virus - Viruses that were detected by the system.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics application-security report view-by violation
```

```
show analytics application-security report view-by violation drilldown { { entity severity values { Error } } }
```

```
send-mail analytics application-security report view-by ip measures {requests} limit 20 order-by { { measure requests sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

requests

The total number of requests for the selected filter (entity).

occurrences

Number of occurrences for the selected filter (relevant for attack-type, violation and ip-address-intelligence entities)

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 analytics application-security report(1)

analytics application-security scheduled-report

NAME

scheduled-report - Configure scheduled reports for application security (ASM).

MODULE

analytics application-security

SYNTAX

Configure the scheduled-report component within the analytics application-security module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the application security module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration

defined in `asm_smtp_conf`.

`modify scheduled-report myScheduledReport smtp-config none`

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to `person@domain.co` using the smtp configuration defined in `asm_smtp_conf`.

`list scheduled-report`

Displays all of the application security scheduled reports.

OPTIONS

`email-addresses`

A list of the email addresses of the recipients that receive the scheduled report.

`first-time`

First scheduled report time. Must be after current time and rounded up to the next round hour.

`frequency`

The scheduled report frequency. Example: `every-6-hours` means that the report will be generated and sent every 6 hours.

`include-total`

Enables or disables including a summary (Overall result) entity in results.

`multi-leveled-report`

Defines a custom multi-leveled report. Mutually exclusive with `predefined-report-name`. The multi-leveled-report definition contains the following parameters:

`chart-path`

A list of entities that define the scope in which the report will be displayed. For example: a chart path `{ violation url }` means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics application-security report.

`limit`

The number of view-by entities displayed in the scheduled report.

`time-diff`

The time range for the report.

`view-by`

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics application-security report.

`measures`

The measures which are available for the selected entities.

`predefined-report-name`

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with `multi-leveled-report`.

`smtp-config`

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

`device-group`

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

`list`, `modify`, `show`, `tmsh`, `analytics application-security report`, `sys smtp-server`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2018-01-16 analytics application-security scheduled-report(1)

analytics asm-bypass report

NAME

report - Displays an asm-bypass analytics report.

MODULE

analytics asm-bypass

SYNTAX

Show, save or send an analytics asm-bypass report using the syntax shown in the following sections.

DISPLAY

show report view-by [slot | memory]

options:

drilldown {

{

entity [slot | memory]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [slot | memory]

options:

drilldown {

{

entity [slot | memory]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [slot | memory]

options:

drilldown {

{

entity [slot | memory]

values

{

[value ...]

}

} ...

}

email-addresses {

[email address ...]

}

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

```
[measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate asm-bypass analytics reports. You can generate an asm-bypass analytics report for the following entities:

• slot - Blade Number

• memory - BD memory

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics asm-bypass report view-by slot
```

```
show analytics asm-bypass report view-by slot drilldown { { entity slot values { slot_index } } }
```

```
end-mail analytics asm-bypass report view-by slot measures { backlog-messages } limit 20 order-by { { measure backlog-messages sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

backlog-messages

The sum of backlog messages for the selected filter (entity).

http-requests

The sum of HTTP requests for the selected filter (entity).

transactions-bypass

The sum of transactions bypass for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2014-11-19 analytics asm-bypass report(1)

analytics asm-bypass scheduled-report

NAME

scheduled-report - Configure scheduled reports for ASM bypass.

MODULE

analytics asm-bypass

SYNTAX

Configure the scheduled-report component within the analytics asm-bypass module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the ASM bypass module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the ASM bypass scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics asm-bypass report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics asm-bypass report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics asm-bypass report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics asm-bypass scheduled-report(1)

analytics asm-cpu report

NAME

report - Displays an asm-cpu analytics report.

MODULE

analytics asm-cpu

SYNTAX

Show, save or send an analytics asm-cpu report using the syntax shown in the following sections.

DISPLAY

show report view-by [virtual | slot]

options:

drilldown {

```

{
entity [ virtual | slot ]
values
{
[value ...]
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]

```

SAVE

```

save report view-by [ virtual | slot ]
options:
drilldown {
{
entity [ virtual | slot ]
values
{
[value ...]
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]

```

SEND

```

send-mail report view-by [ virtual | slot ]
options:
drilldown {
{
entity [ virtual | slot ]
values
{
[value ...]
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate asm-cpu analytics reports. You can generate an asm-cpu analytics report for the following entities:

- virtual - Virtual servers.

Â· slot - Blade Number

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics asm-cpu report view-by slot
```

```
show analytics asm-cpu report view-by slot drilldown { { entity slot values { slot_index } } }
```

```
send-mail analytics asm-cpu report view-by slot measures { bd-cpu-utilization } limit 20 order-by { { measure  
bd-cpu-utilization sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

bd-cpu-utilization

The BD CPU utilization for the selected filter (entity).

tmm-cpu-utilization

The TMM CPU utilization for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2014-11-13 analytics asm-cpu report(1)

analytics asm-cpu scheduled-report

NAME

scheduled-report - Configure scheduled reports for ASM CPU.

MODULE

analytics asm-cpu

SYNTAX

Configure the scheduled-report component within the analytics asm-cpu module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the ASM CPU module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the ASM CPU scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics asm-cpu report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics asm-cpu report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics asm-cpu report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics asm-cpu scheduled-report(1)

analytics asm-enforced-entities report

NAME

report - Displays an ASM Enforced Entities analytics report.

MODULE

analytics asm-enforced-entities

SYNTAX

Show, save or send an analytics asm-enforced-entities report using the syntax shown in the following sections.

DISPLAY

show report view-by [policy]

options:

drilldown {

{

entity [policy]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

```

}
range [date range]

SAVE
save report view-by [ policy ]
options:
  drilldown {
  {
entity [ policy ]
values
{
[value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]

```

```

SEND
send-mail report view-by [ policy ]
options:
  drilldown {
  {
entity [ policy ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate ASM Enforced Entities analytics reports. You can generate an ASM Enforced Entities analytics report for the following entities:

• policy - Security Policy.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics asm-enforced-entities report view-by policy
```

```
show analytics asm-enforced-entities report view-by policy drilldown { { entity policy values { policy_1
policy_2 } } }
```

```
send-mail analytics asm-enforced-entities report view-by policy measures { not-enforced-entities-count } limit
20 order-by { { measure not-enforced-entities-count sort-type desc } } format pdf email-addresses {
some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device
Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

enforced-entities-count

The total number of enforced entities in a specific security policy.

not-enforced-entities-count

The total number of entities that are not enforced in a specific security policy.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics asm-enforced-entities report(1)

analytics asm-learning-suggestions report

NAME

report - Displays an ASM Learning Suggestions analytics report.

MODULE

analytics asm-learning-suggestions

SYNTAX

Show, save or send an analytics asm-learning-suggestions report using the syntax shown in the following sections.

DISPLAY

show report view-by [policy]

options:

drilldown {

{

entity [policy]

```

values
{
  [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

```

SAVE
save report view-by [ policy ]
options:
  drilldown {
    {
  entity [ policy ]
  values
  {
    [value ...]
  } ...
  }
  }
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

```

SEND
send-mail report view-by [ policy ]
options:
  drilldown {
    {
  entity [ policy ]
  values
  {
    [value ...]
  } ...
  }
  }
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate ASM Learning Suggestions analytics reports. You can generate an ASM Learning Suggestions analytics report for the following entities:

• policy - Security Policy.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics asm-learning-suggestions report view-by policy
```

```
show analytics asm-learning-suggestions report view-by policy drilldown { { entity policy values { policy_1 policy_2 } } }
```

```
send-mail analytics asm-learning-suggestions report view-by policy measures { pending-high-score-suggestions-count } limit 20 order-by { { measure pending-high-score-suggestions-count sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

accepted-suggestions-count

The total number of accepted suggestions for a specific security policy.

ignored-suggestions-count

The total number of ignored suggestions for a specific security policy.

pending-high-score-suggestions-count

The total number of pending suggestions with a score of at least 50 for a specific security policy.

pending-low-score-suggestions-count

The total number of pending suggestions with a score less than 50 for a specific security policy.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

analytics asm-memory report

NAME

report - Displays an asm-memory analytics report.

MODULE

analytics asm-memory

SYNTAX

Show, save or send an analytics asm-memory report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ slot ]
options:
  drilldown {
  {
  entity [ slot ]
  values
  {
  [value ...]
  }
  } ...
  }
  field-fmt
  include-total
  include-others
  limit [number of rows]
  measures {
  [measure name ...]
  }
  order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
  }
  range [date range]
```

SAVE

```
save report view-by [ slot ]
options:
  drilldown {
  {
  entity [ slot ]
  values
  {
  [value ...]
  }
  } ...
  }
  file [ file name ]
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
  [measure name ...]
  }
  order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
  }
  range [date range]
```

SEND

```
send-mail report view-by [ slot ]
options:
  drilldown {
  {
  entity [ slot ]
  values
  {
  [value ...]
  }
  } ...
  }
  email-addresses {
  [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
```

```

include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate asm-memory analytics reports. You can generate a ASM memory analytics report for the following entities:

• slot - Slot ID

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics asm-memory report view-by slot
```

```
show analytics asm-memory report view-by slot drilldown { { entity slot values { slot_index } } }
```

```
send-mail analytics asm-memory report view-by slot measures { total-swap-size } limit 20 order-by { { measure total-swap-size sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

bd-memory-utilization

The BD memory utilization out of total assigned to BD (in percents) for the selected filter (entity).

tmm-memory-util

The TMM memory utilization out of total assigned to TMM (in percents) for the selected filter (entity).

bd-swap-size

The swap used by BD in MBs for the selected filter (entity).

total-swap-size

The total swap used in MBs for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2014-11-13 analytics asm-memory report(1)

analytics asm-memory scheduled-report

NAME

scheduled-report - Configure scheduled reports for ASM memory.

MODULE

analytics asm-memory

SYNTAX

Configure the scheduled-report component within the analytics asm-memory module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the ASM memory module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
```

```
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

list scheduled-report

Displays all of the ASM memory scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics asm-memory report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics asm-memory report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics asm-memory report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics asm-memory scheduled-report(1)

analytics asm-policy-changes report

NAME

report - Displays an ASM Policy Changes analytics report.

MODULE

analytics asm-policy-changes

SYNTAX

Show, save or send an analytics asm-policy-changes report using the syntax shown in the following sections.

DISPLAY

show report view-by [policy | user]

options:

```
drilldown {
  {
entity [ policy | user ]
values
{
  [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SAVE

save report view-by [policy | user]

options:

```
drilldown {
  {
entity [ policy | user ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SEND

send-mail report view-by [policy | user]

options:

```
drilldown {
  {
entity [ policy | user ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

smtp-config-override [smtp configuration object name]

DESCRIPTION

Use this command to generate ASM Policy Changes analytics reports. You can generate an ASM Policy Changes analytics report for the following entities:

• policy - Security Policy.

• user - User.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics asm-policy-changes report view-by policy
```

```
show analytics asm-policy-changes report view-by policy drilldown { { entity user values { user_1 user_2 } } }
```

```
send-mail analytics asm-policy-changes report view-by policy measures { policy-changes-count } limit 20 order-by { { measure policy-changes-count sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

policy-changes-count

The total number of changes made to the policy by a specific user.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

analytics asm-violation report

NAME

report - Displays an Application Security Violation analytics report.

MODULE

analytics asm-violation

SYNTAX

Show, save or send an analytics asm-violation report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ application | attack-type | microservice | policy | protocol | request-type | severity | violation |
violation-rating | virtual | virus ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ application | attack-type | microservice | policy | protocol | request-type | severity | violation |
violation-rating | virtual | virus ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
field-fmt
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SAVE

```
save report view-by [ application | attack-type | microservice | policy | protocol | request-type | severity | violation |
violation-rating | virtual | virus ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ application | attack-type | microservice | policy | protocol | request-type | severity | violation |
violation-rating | virtual | virus ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
file [ file name ]
```

```
format [ csv-aggregated | csv-time-series | pdf ]
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SEND

```
send-mail report view-by [ application | attack-type | microservice | policy | protocol | request-type | severity | violation |
violation-rating | virtual | virus ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ application | attack-type | microservice | policy | protocol | request-type | severity | violation |
violation-rating | virtual | virus ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```

} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate Application Security Violation analytics reports. You can generate an Application Security Violation analytics report for the following entities:

- application - Application services.
- attack-type - Attack type of the illegal request.
- microservice - Microservice.
- policy - Security policy.
- protocol - Protocols of requests (HTTP/HTTPS).
- request-type - Enforcement action on request(Legal or Alarmed/Blocked/Dropped).
- severity - Violation severities.
- violation - Violation types.
- violation-rating - Violation ratings.
- virtual - Virtual servers.
- virus - Viruses that were detected by the system.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics asm-violation report view-by application
```

```
show analytics asm-violation report view-by application drilldown { { entity attack-type values {
attack-type_Value } } }
```

```
send-mail analytics asm-violation report view-by application measures { requests } limit 20 order-by { {
measure requests sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

occurrences

Occurrence count.

requests

Request count.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, Itm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics asm-violation report(1)

analytics asm-violation scheduled-report

NAME

scheduled-report - Configure scheduled reports for ASM violation.

MODULE

analytics asm-violation

SYNTAX

Configure the scheduled-report component within the analytics asm-violation module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE

delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the ASM violation module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from being generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com } frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit 5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the ASM violation scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics asm-violation report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics asm-violation report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsr, analytics asm-violation report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

analytics bot-defense-event report

NAME

report - Displays a Bot Defense Event analytics report.

MODULE

analytics bot-defense-event

SYNTAX

Show, save or send a analytics bot-defense-event report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ bot-category | bot-classification | implied-mitigation-action | micro-service-name | micro-service-type |
mobile-app-emulation-mode | mobile-app-jailbroken | mobile-app-name | mobile-app-version |
mobile-is-human-behavior | profile | virtual ]
```

options:

```
drilldown {
```

```
{
entity [ bot-category | bot-classification | implied-mitigation-action | micro-service-name | micro-service-type |
mobile-app-emulation-mode | mobile-app-jailbroken | mobile-app-name | mobile-app-version |
mobile-is-human-behavior | profile | virtual ]
```

values

```
{
[value ...]
```

```
} ...
}
```

field-fmt

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
measure [ measure name ]
sort-type [ asc / desc ]
```

```
} ...
}
```

```
range [date range]
```

SAVE

```
save report view-by [ bot-category | bot-classification | implied-mitigation-action | micro-service-name | micro-service-type |
mobile-app-emulation-mode | mobile-app-jailbroken | mobile-app-name | mobile-app-version |
mobile-is-human-behavior | profile | virtual ]
```

options:

```
drilldown {
```

```
{
entity [ bot-category | bot-classification | implied-mitigation-action | micro-service-name | micro-service-type |
mobile-app-emulation-mode | mobile-app-jailbroken | mobile-app-name | mobile-app-version |
mobile-is-human-behavior | profile | virtual ]
```

values

```
{
[value ...]
```

```
} ...
}
```

```
file [ file name ]
```

```
format [ csv-aggregated | csv-time-series | pdf ]
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
measure [ measure name ]
sort-type [ asc / desc ]
```

```
} ...
}
```

```
range [date range]
```

SEND

```
send-mail report view-by [ bot-category | bot-classification | implied-mitigation-action | micro-service-name | micro-service-type |
```

```

mobile-app-emulation-mode | mobile-app-jailbroken | mobile-app-name | mobile-app-version |
mobile-is-human-behavior | profile | virtual ]
options:
drilldown {
{
entity [ bot-category | bot-classification | implied-mitigation-action | micro-service-name | micro-service-type |
mobile-app-emulation-mode | mobile-app-jailbroken | mobile-app-name | mobile-app-version |
mobile-is-human-behavior | profile | virtual ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate Bot Defense Event analytics reports. You can generate a Bot Defense Event analytics report for the following entities:

- bot-category - Bot Defense Category.
- bot-classification - Bot Classification.
- implied-mitigation-action - Implied Mitigation Action.
- micro-service-name - Bot Defense Micro Service Name.
- micro-service-type - Bot Defense Micro Service Type.
- mobile-app-emulation-mode - Mobile App Emulation Mode.
- mobile-app-jailbroken - Mobile App Jail Broken Rooted.
- mobile-app-name - Mobile App Display Name.
- mobile-app-version - Mobile App Version.
- mobile-is-human-behavior - Mobile Is Human Behavior.
- profile - Bot Defense Profile.
- virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics bot-defense-event report view-by bot-category
```

```
show analytics bot-defense-event report view-by bot-category drilldown { { entity bot-classification values {
some_value } } }
```

```
send-mail analytics bot-defense-event report view-by bot-category measures { occurrences } limit 20 order-by {
{ measure occurrences sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device
Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list
Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown
Specifies specific entities that are used as a filter.

email-addresses
Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format
Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others
Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total
Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit
Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures
Specifies a list of measures that can be used with the chosen entity type. The options are:

occurrences
Occurrence count.

order-by
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics bot-defense-event report(1)

analytics cpu-per-vip report

NAME

report - Displays a CPU Per Virtual analytics report.

MODULE

analytics cpu-per-vip

SYNTAX

Show, save or send a analytics cpu-per-vip report using the syntax shown in the following sections.

DISPLAY

show report view-by [slot | virtual]

options:

drilldown {

{

entity [slot | virtual]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

```

}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SAVE
save report view-by [ slot | virtual ]
options:
  drilldown {
    {
      entity [ slot | virtual ]
      values
      {
        [value ...]
      }
    } ...
  }
  file [ file name ]
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc / desc ]
    } ...
  }
  range [date range]

SEND
send-mail report view-by [ slot | virtual ]
options:
  drilldown {
    {
      entity [ slot | virtual ]
      values
      {
        [value ...]
      }
    } ...
  }
  email-addresses {
    [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc / desc ]
    } ...
  }
  range [date range]
  smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate CPU Per Virtual analytics reports. You can generate a CPU Per Virtual analytics report for the following entities:

• slot - Slot ID.

• virtual - Virtual server.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics cpu-per-vip report view-by slot
```

```
show analytics cpu-per-vip report view-by slot drilldown { { entity virtual values { virtual_1 virtual_2 } } }
```

```
send-mail analytics cpu-per-vip report view-by slot measures { avg-1min } limit 20 order-by { { measure avg-1min sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-1min

Average one minute.

avg-5min

Average five minute.

avg-5sec

Average five seconds.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics cpu-per-vip report(1)

analytics cpu report

NAME

report - Displays an cpu analytics report.

MODULE

analytics cpu

SYNTAX

Show, save or send an analytics cpu report using the syntax shown in the following sections.

DISPLAY

show report view-by [slot | cpu]

options:

```
drilldown {
  {
  entity [ slot | cpu ]
  values
  {
  [value ...]
  }
  } ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SAVE

save report view-by [slot | cpu]

options:

```
drilldown {
  {
  entity [ slot | cpu ]
  values
  {
  [value ...]
  }
  } ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SEND

send-mail report view-by [slot | cpu]

options:

```
drilldown {
  {
  entity [ slot | cpu ]
  values
  {
  [value ...]
  }
  } ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate cpu analytics reports. You can generate a IP-layer analytics report for the following entities:

• slot - Slot ID

• cpu - CPU number ID

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics cpu report view-by cpu
```

```
show analytics cpu report view-by cpu drilldown { { entity cpu values { 1 } } }
```

```
send-mail analytics cpu report view-by slot measures { cpu-usage } limit 20 order-by { { measure cpu-usage  
sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

cpu-usage

The average time percentage of real CPU usage (user + system + nice) usage for the selected filter (entity).

io The average time percentage of (iowait + irq + softirq) for the selected filter (entity).

stolen

The average time percentage of virtual CPU waits for a real CPU while the hypervisor is servicing another virtual processor.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

analytics cpu scheduled-report

NAME

scheduled-report - Configure scheduled reports for CPU.

MODULE

analytics cpu

SYNTAX

Configure the scheduled-report component within the analytics cpu module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the CPU module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the CPU scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics cpu report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics cpu report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics cpu report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics cpu scheduled-report(1)

analytics device-traffic report

NAME

report - Displays a Device Traffic Stats analytics report.

MODULE

analytics device-traffic

SYNTAX

Show, save or send a analytics device-traffic report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ cpu-num | slot-id ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ cpu-num | slot-id ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
field-fmt
```

```

include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

```

SAVE
save report view-by [ cpu-num | slot-id ]
options:
  drilldown {
    {
      entity [ cpu-num | slot-id ]
      values
      {
        [value ...]
      }
    } ...
  }
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

```

SEND
send-mail report view-by [ cpu-num | slot-id ]
options:
  drilldown {
    {
      entity [ cpu-num | slot-id ]
      values
      {
        [value ...]
      }
    } ...
  }
  email-addresses {
    [email address ...]
  }
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate Device Traffic Stats analytics reports. You can generate a Device Traffic Stats analytics report for the following entities:

Â· cpu-num - CPU num.

Â· slot-id - Slot id.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics device-traffic report view-by cpu-num
```

```
show analytics device-traffic report view-by cpu-num drilldown { { entity slot-id values { 0 } } }
```

```
send-mail analytics device-traffic report view-by cpu-num measures { client-bits-out } limit 20 order-by { {
measure client-bits-out sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-client-concurrent-connections

Avg Client Concurrent connections.

avg-server-concurrent-connections

Avg Server Concurrent connections.

client-bits-in

Client Incoming bits.

client-bits-out

Client Outgoing bits.

client-connections

Client Connections.

client-packets-in

Client Incoming packets.

client-packets-out

Client Outgoing packets.

max-client-concurrent-connections

Maximum Client Concurrent connections at peak.

max-server-concurrent-connections

Maximum Server Concurrent connections at peak.

requests

Requests.

server-bits-in

Server Incoming bits.

server-bits-out

Server Outgoing bits.

server-connections

Server Connections.

server-packets-in

Server Incoming packets.

server-packets-out

Server Outgoing packets.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics device-traffic report(1)

analytics device-traffic scheduled-report

NAME

scheduled-report - Configure scheduled reports for device traffic.

MODULE

analytics device-traffic

SYNTAX

Configure the scheduled-report component within the analytics device-traffic module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE

delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the device traffic module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

list scheduled-report

Displays all of the device traffic scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics device-traffic report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics device-traffic report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics device-traffic report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics device-traffic scheduled-report(1)

analytics disk-info report

NAME

report - Displays an disk-info analytics report.

MODULE

analytics disk-info

SYNTAX

Show, save or send an analytics disk-info report using the syntax shown in the following sections.

DISPLAY

show report view-by [slot]

options:

drilldown {

{

entity [slot]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [slot]

options:

drilldown {

{

entity [slot]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [slot]

options:

drilldown {

{

entity [slot]

values

{

[value ...]

}

} ...

}

email-addresses {

[email address ...]

}

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

```
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate disk-info analytics reports. You can generate a disk-info analytics report for the following entities:

• slot - Slot ID

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics disk-info report view-by slot
```

```
show analytics disk-info report view-by slot drilldown { { entity slot values { 7 } } }
```

```
send-mail analytics disk-info report view-by slot measures { total-ios } limit 20 order-by { { measure total-ios sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

total-ios

The total number of I/O for the selected filter (entity).

read-operations

The total number of read operations for the selected filter (entity).

read-merged

The total number of merged reads for the selected filter (entity).

write-operations

The total number of write operations for the selected filter (entity).

write-merged

The total number of merged writes for the selected filter (entity).

read-bytes

The total number of read bytes for the selected filter (entity).

write-bytes

The total number of write bytes for the selected filter (entity).

max-read-latency

The max value of read latency for the selected filter (entity).

average-read-latency

The average number of average read latency for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, itm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-07-10 analytics disk-info report(1)

analytics disk-info scheduled-report

NAME

scheduled-report - Configure scheduled reports for disk information.

MODULE

analytics disk-info

SYNTAX

Configure the scheduled-report component within the analytics disk-info module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the disk information module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

list scheduled-report

Displays all of the disk information scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics disk-info report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics disk-info report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics disk-info report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics disk-info scheduled-report(1)

NAME

report - Displays a DNS Cache Resolver analytics report.

MODULE

analytics dns-cache-resolver

SYNTAX

Show, save or send a analytics dns-cache-resolver report using the syntax shown in the following sections.

DISPLAY

show report view-by [name]

options:

drilldown {

{

entity [name]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [name]

options:

drilldown {

{

entity [name]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [name]

options:

drilldown {

{

entity [name]

values

{

[value ...]

}

} ...

}

email-addresses {

[email address ...]

}

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

```
{
  measure [ measure name ]
  sort-type [ asc / desc ]
} ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate DNS Cache Resolver analytics reports. You can generate a DNS Cache Resolver analytics report for the following entities:

• name - DNS cache resolver name.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics dns-cache-resolver report view-by name
```

```
show analytics dns-cache-resolver report view-by name drilldown { { entity name values { name_1 } } }
```

```
send-mail analytics dns-cache-resolver report view-by name measures { rpz-rewrite } limit 20 order-by { {
measure rpz-rewrite sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

hit-ratio

DNS cache resolver hit ratio.

rpz-rewrite

DNS cache resolver rpz rewrites.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

analytics dns-profile report

NAME

report - Displays a DNS Profiles analytics report.

MODULE

analytics dns-profile

SYNTAX

Show, save or send a analytics dns-profile report using the syntax shown in the following sections.

DISPLAY

show report view-by [name | vs-name]

options:

drilldown {

{

entity [name | vs-name]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [name | vs-name]

options:

drilldown {

{

entity [name | vs-name]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [name | vs-name]

options:

drilldown {

{

entity [name | vs-name]

values

{

```

[value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate DNS Profiles analytics reports. You can generate a DNS Profiles analytics report for the following entities:

• name - name.

• vs-name - vs names.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics dns-profile report view-by name
```

```
show analytics dns-profile report view-by name drilldown { { entity vs-name values { vs_1 vs_2 } } }
```

```
send-mail analytics dns-profile report view-by name measures { average-responses } limit 20 order-by { {
measure average-responses sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

average-a-reqs

IPV4 address requests per seconds.

average-aaaa-reqs

IPV6 address requests per seconds.

average-any-reqs

Retrieves all the available type for given name per seconds.

average-cname-reqs
Alias of one name to another: the DNS lookup will continue by retrying the lookup with the new name per seconds.

average-dns64fails
IPv6 To IPv4 Requests per seconds.

average-dns64reqs
IPv6 To IPv4 Requests per seconds.

average-dns64trans
IPv6 To IPv4 Transfers per seconds.

average-drops
DNS drops per seconds.

average-fast-dns-allowed
Last action allowed per seconds.

average-fast-dns-drops
Last action drops per seconds.

average-fast-dns-queries
Rapid Responses fast queries per seconds.

average-fast-dns-resp-ne
Last Action no error per seconds.

average-fast-dns-resp-nx
Last action NX domain per seconds.

average-fast-dns-resp-rf
Last action refused per seconds.

average-fast-dns-resp-tc
last action truncated per seconds.

average-fast-dns-responses
Rapid responses per seconds.

average-hdr-aa
Authoritative Answers per seconds.

average-hdr-ad
Authenticated data per seconds.

average-hdr-ra
Recursion available per seconds.

average-hdr-tc
Truncated per seconds.

average-hints
Cache of the DNS servers list per seconds.

average-hit-ratio
DNS Hit Ratio - Responses / (Responses + Requests) per seconds.

average-hw-inspected
Inspected per seconds.

average-improperly-formatted
Hardware handled packets improperly formatted per seconds.

average-invalid-authority
Invalid Authority per seconds.

average-malformed
Protocol violations improperly formatted per seconds.

average-mx-reqs
Maps a domain name to a list of message transfer agents for that domain per seconds.

average-naptr-reqs
NAPTR is a type of resource record in the DNS per seconds.

average-ns-reqs
Delegates a DNS zone to use the given authoritative name servers per seconds.

average-other-reqs
Other Requests per seconds.

average-ptr-reqs
Pointer records are used to map a network interface (IP) to a host name per seconds.

average-queries
DNS queries per seconds.

average-r-noerror
No Error per seconds.

average-r-nxdomain
Domain per seconds.

average-r-refused
Query refused per seconds.

average-r-servfail
Server Failure per seconds.

average-rejects
rejects per seconds.

average-responses
DNS responses per seconds.

average-soa-reqs
Specifies authoritative information about a DNS zone per seconds.

average-srv-reqs
A Service record is a specification of data in the Domain Name System defining the location per seconds.

average-success-queries
Query completed successfully per seconds.

average-todns
DNS todns per seconds.

average-txt-reqs
Originally for arbitrary human-readable text in a DNS record per seconds.

order-by
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-05 analytics dns-profile report(1)

analytics dns-protocol scheduled-report

NAME

scheduled-report - Configure scheduled reports for DNS protocol.

MODULE

analytics dns-protocol

SYNTAX

Configure the scheduled-report component within the analytics dns-protocol module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

```
frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
include-total [enabled | disabled]
multi-leveled-report {
chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
limit [number of rows]
time-diff [last-hour | last-day | last-week | last-month | last-year]
view-by { entity name [string] }
measures [none | add | delete | modify | replace-all-with] { measure name [string] }
}
predefined-report-name [name]
smtp-config [name]
device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the DNS protocol module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the DNS protocol scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics dns-protocol report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics dns-protocol report.

measures

The measures which are available for the selected entities.

predefined-report-name
Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config
Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group
Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics dns-protocol report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics dns-protocol scheduled-report(1)

analytics dns-rpz report

NAME

report - Displays a DNS Response Policy Zones analytics report.

MODULE

analytics dns-rpz

SYNTAX

Show, save or send a analytics dns-rpz report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ rpz ]
options:
  drilldown {
  {
  entity [ rpz ]
  values
  {
  [value ...]
  }
  } ...
  }
  field-fmt
  include-total
  include-others
  limit [number of rows]
  measures {
  [measure name ...]
  }
  order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
  }
  range [date range]
```

SAVE

```
save report view-by [ rpz ]
options:
  drilldown {
  {
  entity [ rpz ]
  values
  {
  [value ...]
  }
  } ...
  }
  file [ file name ]
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
```

```

measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ rpz ]
options:
  drilldown {
    {
  entity [ rpz ]
  values
  {
    [value ...]
  } ...
}
  email-addresses {
    [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc / desc ]
    } ...
  }
  range [date range]
  smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate DNS Response Policy Zones analytics reports. You can generate a DNS Response Policy Zones analytics report for the following entities:

- rpz - Response policy zone.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics dns-rpz report view-by rpz
```

```
show analytics dns-rpz report view-by rpz drilldown { { entity rpz values { rpz_name } } }
```

```
send-mail analytics dns-rpz report view-by rpz measures { rpz-queries } limit 20 order-by { { measure rpz-queries sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

queries

Queries.

rpz-queries

Response policy zone queries.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, Itm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics dns-rpz report(1)

analytics dns report

NAME

report - Displays a DNS analytics report.

MODULE

analytics dns

SYNTAX

Show, save or send an analytics dns report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ activity-type | application | attack-id | client-ip | country | country-code | dns-transaction-outcome | domain-name | dos-
options:
drilldown {
  {
entity [ activity-type | application | attack-id | client-ip | country | country-code | dns-transaction-outcome | domain-name | dos-profile | mitig
values
{
[value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc | desc ]
} ...
}
range [date range]
```

SAVE

```
save report view-by [ activity-type | application | attack-id | client-ip | country | country-code | dns-transaction-outcome | domain-name | dos-
```

```

options:
  drilldown {
  {
entity [ activity-type | application | attack-id | client-ip | country | country-code | dns-transaction-outcome | domain-name | dos-profile | mitig
values
{
[value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc | desc ]
} ...
}
}
range [date range]

```

SEND

```

send-mail report view-by [ activity-type | application | attack-id | client-ip | country | country-code | dns-transaction-outcome | domain-name |
options:
  drilldown {
  {
entity [ activity-type | application | attack-id | client-ip | country | country-code | dns-transaction-outcome | domain-name | dos-profile | mitig
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc | desc ]
} ...
}
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate DNS analytics reports. You can generate a DNS analytics report for the following entities:

- Â· activity-type - Activity type.
- Â· application - Application services (iApps(tm)).
- Â· attack-id - (Only available with AFM is provisioned) DoS Attack ID.
- Â· client-ip - DNS query source/client IP address.
- Â· country - Country.
- Â· country-code - Country code.
- Â· dns-transaction-outcome - Request outcome.
- Â· domain-name - Queried domain name.
- Â· dos-profile - DoS profile.
- Â· mitigation - Mitigation.
- Â· query-type - DNS query type.
- Â· suspected-ip - Suspected address IP.
- Â· trigger - Trigger.

Â· vector - Attack vector.

Â· virtual - Virtual server.

EXAMPLES

```
show analytics dns report view-by virtual
```

```
show analytics dns report view-by query-type drilldown { { entity virtual values { /Common/v1 } } }
```

```
send-mail analytics dns report view-by client-ip limit 20 format pdf email-addresses {  
some.one@someaddress.com }
```

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

attacks-count

The total number of DNS attacks for the specified view-by entity.

packets

The total number of DNS packets for the specified view-by entity.

packets-per-second

The average number of DNS packets for the specified view-by entity.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile dns, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

analytics dns scheduled-report

NAME

scheduled-report - Configure scheduled reports for DNS.

MODULE

analytics dns

SYNTAX

Configure the scheduled-report component within the analytics dns module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the DNS module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the DNS scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics dns report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics dns report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics dns report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics dns scheduled-report(1)

analytics dos-l3 report

NAME

report - Displays a DoS (Layers 3-4) prevention analytics report.

MODULE

analytics dos-l3

SYNTAX

Show, save or send an analytics dos-l3 report using the syntax shown in the following sections.

DISPLAY

show report view-by [action | activity-type | application | attack-id | category | client-ip | country | country-code | dest-country | dest-country-dos-profile | mitigation | server-ip | suspected-ip | trigger | vector | virtual | vlan | vlan-group]

options:

drilldown {

{

entity [action | activity-type | application | attack-id | category | client-ip | country | country-code | dest-country | dest-country-code | dos-profile | mitigation | server-ip | suspected-ip | trigger | vector | virtual | vlan | vlan-group]

values

{
[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

```

order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]

```

SAVE

```

save report view-by [ action | activity-type | application | attack-id | category | client-ip | country | country-code | dest-country | dest-country-code | dos-profile | mitigation | server-ip | suspected-ip | trigger | vector | virtual | vlan | vlan-group ]
options:
  drilldown {
    {
entity [ action | activity-type | application | attack-id | category | client-ip | country | country-code | dest-country | dest-country-code | dos-profile | mitigation | server-ip | suspected-ip | trigger | vector | virtual | vlan | vlan-group ]
values
{
  [value ...]
}
} ...
}
  file [ file name ]
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc | desc ]
    } ...
  }
  range [date range]

```

SEND

```

send-mail report view-by [ action | activity-type | application | attack-id | category | client-ip | country | country-code | dest-country | dest-country-code | dos-profile | mitigation | server-ip | suspected-ip | trigger | vector | virtual | vlan | vlan-group ]
options:
  drilldown {
    {
entity [ action | activity-type | application | attack-id | category | client-ip | country | country-code | dest-country | dest-country-code | dos-profile | mitigation | server-ip | suspected-ip | trigger | vector | virtual | vlan | vlan-group ]
values
{
  [value ...]
}
} ...
}
  email-addresses {
    [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc | desc ]
    } ...
  }
  range [date range]
  smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate DoS (Layers 3-4) prevention analytics reports. You can generate a DoS prevention analytics report for the following entities:

Â· action - Action taken (allowed/dropped).

Â· activity-type - Activity type.

Â· application - Application services (iApps(tm)).

Â· attack-id - DoS attack ID.

Â· category - Attack category.

Â· client-ip - Source/client IP address.

- Â· country - Country.
- Â· country-code - Country code.
- Â· dest-country - Destination country.
- Â· dest-country-code - Destination country code.
- Â· dos-profile - DoS profile.
- Â· mitigation - Mitigation.
- Â· server-ip - Server address IP.
- Â· suspected-ip - Suspect address IP.
- Â· trigger - Trigger.
- Â· vector - Attack vector.
- Â· virtual - Virtual server.
- Â· vlan - VLAN.
- Â· vlan-group - VLAN Group.

EXAMPLES

show analytics dos-l3 report view-by virtual

show analytics dos-l3 report view-by attack-type drilldown { { entity virtual values { /Common/v1 } } }

send-mail analytics dos-l3 report view-by source-ip limit 20 format pdf email-addresses { some.one@someaddress.com }

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

allowed-requests

The total number of packets that were received by the virtual server(/s)s

allowed-requests-per-second

The average number of packets that were received by the virtual server(/s)s

attacks-count

The total number of attacks for the selected view-by entity.

dropped-requests

The total number of packets that were dropped by the virtual server(/s)s

dropped-requests-per-second

The average number of packets that were dropped by the virtual server(/s)s

total-requests

The total number of packets that were received or dropped by the virtual server(/s)

total-requests-per-second

The average number of packets that were received or dropped by the virtual server(/s)

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2018-02-01 analytics dos-l3 report(1)

analytics dos-l3 scheduled-report

NAME

scheduled-report - Configure scheduled reports for DoS L3.

MODULE

analytics dos-l3

SYNTAX

Configure the scheduled-report component within the analytics dos-l3 module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE

delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the DoS L3 module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF

showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the DoS L3 scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics dos-l3 report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics dos-l3 report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics dos-l3 report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics dos-l3 scheduled-report(1)

analytics dos-l7 report

NAME

report - Displays an HTTP/L7-DoS analytics report.

MODULE

analytics dos-l7

SYNTAX

Show, save or send an analytics dos-l7 report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
options:
drilldown {
  {
entity [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
values
{
  [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
```

SAVE

```
save report view-by [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
options:
drilldown {
  {
entity [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
```

SEND

```
send-mail report view-by [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
options:
drilldown {
  {
entity [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
values
{
```

```

[value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate HTTP analytics reports. You can generate an HTTP analytics report for the following entities:

- Â· activity-type - Activity type.
- Â· application - Application services.
- Â· attack-id - Application/L7 DoS Attack ID.
- Â· behavioral-signature - Behavioral signature.
- Â· browser - Browser.
- Â· client-ip - A single client identified by an IP address.
- Â· client-subnet - Client subnet.
- Â· country - A country from which HTTP/HTTPS traffic was sent to each of the virtual servers.
- Â· country-code - Country code from which HTTP/HTTPS traffic was sent to each of the virtual servers.
- Â· device-id - Device ID.
- Â· dos-profile - DoS Profile.
- Â· http-method - Method.
- Â· http-transaction-outcome - HTTP Transaction outcomes (Blocked/Dropped/Passthrough/etc.)
- Â· mitigation - Mitigation.
- Â· os - OS name.
- Â· pool-member - Pool members.
- Â· response-code - An HTTP response code that was sent back to the client.
- Â· suspected-ip - Suspected address IP.
- Â· trigger - Trigger.
- Â· url - A URL accessed by HTTP or HTTPS.
- Â· user-agent - A browser identifier sent by the client's browser as part of the request for URL.
- Â· vector - Attack vector.
- Â· virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics dos-l7 report view-by virtual measures {average-tps} limit 20
```

Gets the average tps of 20 virtual servers (unordered).

```
show analytics dos-l7 report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps
sort-type desc } }
```

Gets the average tps of the top 20 virtual servers.

```
show analytics dos-l7 report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps
sort-type desc } } range now-3d--now
```

Gets the average tps of the top 20 virtual servers from the last three days.

```
show analytics dos-l7 report view-by virtual drilldown { { entity application values { app } } { entity pool-member values { p1 p2 } } } range now-4d--now-2d measures {average-tps} limit 10 order-by { { measure average-tps sort-type DESC } }
```

Gets the average tps of the top 10 virtual servers (ordered by average tps) on app iApp (out of several monitored) on pool members p1 and p2 (out of five monitored p1-p5) in the interval ranging from two to four days ago.

```
show analytics dos-l7 report view-by response-code drilldown { { entity virtual values { v1 } } } measures { transactions }
```

Gets a distribution of requests per response code on virtual v1.

```
show analytics dos-l7 report view-by country drilldown { { entity application values { app } } } measures { average-concurrent-sessions average-sessions } order-by { { measure average-sessions sort-type DESC } } limit 5
```

Gets the new sessions and average concurrent sessions of the top five countries, ordered by the average concurrent sessions on the application app.

```
show analytics dos-l7 report view-by client-ip drilldown { { entity virtual values { v1 } } } measures { max-page-load-time } limit 1
```

Gets the client IP address with the worst page load time.

```
show analytics dos-l7 report view-by application drilldown { { entity pool-member values { p1 p2 } } } measures { transactions } order-by { { measure transactions } } range now-7d--now
```

Gets the distribution of requests per application on pool members p1 and p2 ordered by the number of requests during the last week.

```
save analytics dos-l7 report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format pdf file report.pdf
```

Gets the average tps of the top 20 virtual servers and exports to a PDF file on the BIG-IP system.

```
save analytics dos-l7 report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format csv-aggregated file report.csv
```

Gets the average tps of the top 20 virtual servers and exports to a CSV file on the BIG-IP system.

```
save analytics dos-l7 report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format csv-time-series file report.csv
```

Gets the average tps over time of the top 10 virtual servers and exports to a CSV file on the BIG-IP system.

```
send-mail analytics dos-l7 report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

Gets the average tps over time of the top 10 virtual servers and sends out an email containing the report as a PDF.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The default value is transactions. The options are:

average-concurrent-sessions

The average number of concurrent sessions for each entity.

average-new-sessions

The average number of new sessions for each entity.

average-page-load-time

The average client page load time for each entity.

average-request-throughput

The average request throughput for each entity.

average-response-throughput

The average response throughput for each entity.

average-server-latency

The average server latency for each entity.

average-tps

The average number of transactions per second for each entity.

client-side-sampled-transactions

The number of transactions sampled for client side page load time.

max-page-load-time

The maximum client page load time for each entity.

max-request-throughput

The maximum request throughput for each entity.

max-response-throughput

The maximum response throughput for each entity.

max-server-latency

The maximum server latency for each entity.

max-tps

The maximum number of transactions per second for each entity.

transactions

The absolute number of transactions for each entity.

min-server-latency

The minimum server latency for each entity.

average-request-size

The average request size for each entity.

average-response-size

The average response size for each entity.

average-application-response-time

The average application response time for each entity.

min-application-response-time

The minimum application response time for each entity.

max-application-response-time

The maximum application response time for each entity.

average-client-ttfb

The average client TTFB for each entity.

min-client-ttfb

The minimum client TTFB for each entity.

max-client-ttfb

The maximum client TTFB for each entity.

average-clientside-network-latency

The average client-side network latency for each entity.

min-clientside-network-latency

The minimum client-side network latency for each entity.

max-clientside-network-latency

The maximum client-side network latency for each entity.

average-serverside-network-latency

The average server-side network latency for each entity.

min-serverside-network-latency

The minimum server-side network latency for each entity.

`max-serverside-network-latency`
The maximum server-side network latency for each entity.

`average-request-duration`
The average request duration for each entity.

`min-request-duration`
The minimum request duration for each entity.

`max-request-duration`
The maximum request duration for each entity.

`average-response-duration`
The average response duration for each entity.

`min-response-duration`
The minimum response duration for each entity.

`max-response-duration`
The maximum response duration for each entity.

`attacks-count`
The total number of attack for each entity.

`valid`
The total number of valid transactions for each entity.

`average-valid-tps`
The average number of valid transactions for each entity.

`mitigated`
The total number of mitigated transaction for each entity.

`average-mitigated-tps`
The average number of mitigated transaction for each entity.

`blocked`
The total number of blocked transactions for each entity.

`average-blocked-tps`
The average number of blocked transactions for each entity.

`incomplete`
The total number of incomplete transactions for each entity.

`average-incomplete-tps`
The average number of incomplete transactions for each entity.

`order-by`
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The default value for measures is previously chosen measures. The default value for sort type is desc (descending).

`range`
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

`smtp-config-override`
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

`show`, `save`, `send-mail`, `tms`, `itm` profile analytics, security dos profile, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2018-07-05 analytics dos-l7 report(1)

analytics dos-vis-attacks report

NAME

report - Displays DoS Visibility Attacks analytics report.

MODULE

analytics dos-vis-attacks

SYNTAX

Show, save or send analytics dos-vis-attacks report using the syntax shown in the following sections.

DISPLAY

show report view-by [attack-id | event-idx | protocol | virtual | dos-profile | trigger | vector | mitigation]

options:

```
drilldown {
  {
    entity [ attack-id | event-idx | protocol | virtual | dos-profile | trigger | vector | mitigation ]
  }
  values
  {
    [value ...]
  }
  } ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]
```

SAVE

save report view-by [attack-id | event-idx | protocol | virtual | dos-profile | trigger | vector | mitigation]

options:

```
drilldown {
  {
    entity [ attack-id | event-idx | protocol | virtual | dos-profile | trigger | vector | mitigation ]
  }
  values
  {
    [value ...]
  }
  } ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]
```

SEND

send-mail report view-by [attack-id | event-idx | protocol | virtual | dos-profile | trigger | vector | mitigation]

options:

```
drilldown {
  {
    entity [ attack-id | event-idx | protocol | virtual | dos-profile | trigger | vector | mitigation ]
  }
  values
  {
    [value ...]
  }
  } ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
```

```
} ...  
}  
range [date range]  
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate DoS Attacks analytics reports. You can generate a DoS Attacks analytics report for the following entities:

- attack-id - Unique attack ID.
- event-idx - Index of an event in an attack.
- protocol - Attacked protocol (HTTP, DNS, SIP, Network).
- virtual - Name of the attacked Virtual Server.
- dos-profile - DoS profile that detects and mitigates the attack.
- trigger - Threshold that was hit to indicate the attack.
- vector - Attack vector (used by AFM, ASM has only 1 vector).
- mitigation - Mitigation to the attack.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics dos-vis-attacks report view-by virtual
```

```
show analytics dos-vis-attacks report view-by virtual drilldown { { entity protocol values { DNS } } }
```

```
send-mail analytics dos-vis-attacks report view-by virtual measures { attack-severity } limit 15 order-by { {  
measure attack-severity sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

event-start

The start moment of the event.

event-end

The end moment of the event.

total-blocked

Total number of blocked transactions.

concurrent-ips

Maximum number of simultaneous attacking IPs per 10 second interval.

severity

Calculated severity of the attack.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-01-20 analytics dos-vis-attacks report(1)

analytics dos-vis-common report

NAME

report - Displays a DoS Common analytics report.

MODULE

analytics dos-vis-common

SYNTAX

Show, save or send a analytics dos-vis-common report using the syntax shown in the following sections.

DISPLAY

show report view-by [activity-type | application | attack-id | client-ip | country | country-code | dos-profile | mitigation | protocol | suspected-ip | trigger | vector | virtual]

options:

drilldown {

{

entity [activity-type | application | attack-id | client-ip | country | country-code | dos-profile | mitigation | protocol | suspected-ip | trigger | vector | virtual]

values

{ [value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [activity-type | application | attack-id | client-ip | country | country-code | dos-profile | mitigation | protocol | suspected-ip | trigger | vector | virtual]

options:

drilldown {

{

entity [activity-type | application | attack-id | client-ip | country | country-code | dos-profile | mitigation | protocol | suspected-ip | trigger | vector | virtual]

values

{ [value ...]

}

} ...

}

```

file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ activity-type | application | attack-id | client-ip | country | country-code | dos-profile | mitigation |
  protocol | suspected-ip | trigger | vector | virtual ]
options:
drilldown {
  {
entity [ activity-type | application | attack-id | client-ip | country | country-code | dos-profile | mitigation |
  protocol | suspected-ip | trigger | vector | virtual ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate DoS Common analytics reports. You can generate a DoS Common analytics report for the following entities:

- Â· activity-type - Tells whether a transaction was created by client regular activity or due to BIG-IP internal activity/injected JavaScripts.
- Â· application - Application services.
- Â· attack-id - Attack's unique ID.
- Â· client-ip - A single client identified by an IP address.
- Â· country - The name of the country from which the traffic arrived.
- Â· country-code - An ISO 3166-1 Alpha-2 country code from which the traffic arrived.
- Â· dos-profile - Name of the DoS Profile involved in classifying relevant traffic as attack.
- Â· mitigation - The mitigation of the attack.
- Â· protocol - The protocol that was attacked (HTTP/SIP/DNS/L3).
- Â· suspected-ip - Is this IP suspected by dos module as "attacking".
- Â· trigger - The trigger of the attack.
- Â· vector - The vector of the attack.
- Â· virtual - Name of the virtual server.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics dos-vis-common report view-by activity-type
```

```
show analytics dos-vis-common report view-by activity-type drilldown { { entity virtual values { virtual_1
virtual_2 } } }
```

```
send-mail analytics dos-vis-common report view-by activity-type measures { network-dropped-requests-per-second } limit 20 order-by { { measure network-dropped-requests-per-second sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

attacks-count

Number of distinct attacks.

average-tps

Average number of transactions per second (tps).

dns-hits-count

The total number of DNS packets.

http-transactions

The absolute number of transactions for each entity.

network-allowed-requests

Total number of attacking network requests that were allowed by AFM.

network-allowed-requests-per-second

Average number of network-attacking requests allowed per second.

network-dropped-requests

Total number of dropped network requests.

network-dropped-requests-per-second

Average number of network-attacking requests dropped per second.

network-total-requests

Total number of attacking requests.

network-total-requests-per-second

Average number of attacking requests per second (allowed and dropped).

packets-per-second

Average DNS packets per second.

sip-hits-count

The total number of SIP requests.

sip-requests-per-sec

Average number of SIP requests per second.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-05 analytics dos-vis-common report(1)

analytics dos-vis-vips report

NAME

report - Displays a DoS Virtual Servers analytics report.

MODULE

analytics dos-vis-vips

SYNTAX

Show, save or send a analytics dos-vis-vips report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ virtual ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ virtual ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
field-fmt
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SAVE

```
save report view-by [ virtual ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ virtual ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
file [ file name ]
```

```
format [ csv-aggregated | csv-time-series | pdf ]
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```

    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ virtual ]
options:
drilldown {
  {
entity [ virtual ]
values
{
[value ...]
}
} ...
}
email-addresses {
[ email address ... ]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
  {
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate DoS Virtual Servers analytics reports. You can generate a DoS Virtual Servers analytics report for the following entities:

• virtual - Name of the virtual server.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics dos-vis-vips report view-by virtual
```

```
show analytics dos-vis-vips report view-by virtual drilldown { { entity virtual values { virtual_1 virtual_2 } } }
```

```
send-mail analytics dos-vis-vips report view-by virtual measures { avg-client-side-concurrent-conns } limit 20
order-by { { measure avg-client-side-concurrent-conns sort-type desc } } format pdf email-addresses {
some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not

including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-abandoned-conns

Average number of abandoned connections per second.

avg-client-in-pps

Average packets rate in from clients per second.

avg-client-in-throughput

Average bits rate in from clients per second.

avg-client-out-pps

Average packets rate out to clients per second.

avg-client-out-throughput

Average bits rate outs to clients per second.

avg-client-side-concurrent-conns

Average client side concurrent connections.

avg-expired-conns

Average number of expired connections per second.

avg-failed-conns

Average number of failed connections per second.

avg-new-client-conns

Average number of new client side connections per second.

avg-new-server-conns

Average number of new server side connections per second.

avg-server-in-pps

Average packets rate in from servers per second.

avg-server-in-throughput

Average bits rate in from servers per second.

avg-server-latency

Average server latency (in ms).

avg-server-latency-health

Health score of the Latency parameter.

avg-server-out-pps

Average packets rate out to servers per second.

avg-server-out-throughput

Average bits rate out to servers per second.

avg-server-side-concurrent-conns

Average server side concurrent connections.

avg-throughput-health

Health score of the Throughput parameter.

concurrent-attacking-ips

The number of monitored simultaneous IPs participating in attacks per 10 second interval.

concurrent-blocked-ips

The number of monitored simultaneous blocked IPs per 10 second interval.

concurrent-conns-health

Health score of the Connections parameter.

concurrent-ips

The number of monitored simultaneous IP connections per 10 second interval.

max-concurrent-attacks

The number of monitored simultaneous attacks per 10 second interval.

max-concurrent-attacks-overtime

The at least number of simultaneous attacks per 10 second interval.

max-concurrent-client-conns

Highest number of simultaneous client side connections observed.

max-concurrent-server-conns

Highest number of simultaneous server side connections observed.

special-concurrent-ips-for-all-vips

The least number of simultaneous IPs connections per 10 second interval for all VIPs.

total-abandoned-conns

Total number of abandoned connections.

total-client-in-bytes
Total bytes in from clients.

total-client-out-bytes
The total client out bytes.

total-expired-conns
Total number of expired connections.

total-failed-conns
Total number of failed connections.

total-health
Overall health score of the virtual server (in %).

total-new-client-conns
Total number of new client side connections.

total-new-server-conns
Total number of new server side connections.

total-server-in-bytes
Total bytes in from servers.

total-server-out-bytes
Total bytes out to servers.

order-by
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics dos-vis-vips report(1)

analytics fw-nat report

NAME

report - Displays a firewall NAT analytics report.

MODULE

analytics fw-nat

SYNTAX

Show, save or send an analytics fw-nat report using the syntax shown in the following sections.

DISPLAY

show report view-by [name]

options:

drilldown {

{

entity [name]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

```

include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]

SAVE
save report view-by [ name ]
options:
  drilldown {
    {
      entity [ name ]
      values
      {
        [value ...]
      } ...
    }
  }
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ name ]
options:
  drilldown {
    {
      entity [ name ]
      values
      {
        [value ...]
      } ...
    }
  }
  email-addresses {
    [email address ...]
  }
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate firewall NAT analytics reports. You can generate a firewall NAT analytics report for the following entities:

• name - Name.

EXAMPLES

```
show analytics fw-nat report view-by name
```

```
show analytics fw-nat report view-by name drilldown { { entity name values { my_name } } }
```

```
send-mail analytics fw-nat report view-by name limit 20 format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. This option must be used with the drilldown option. You can also use it along with include-others.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-source-translation-requests

The average number of FW NAT source translation requests for the selected filter (entity).

avg-source-translation-request-failures

The average number of FW NAT source translation request failures for the selected filter (entity).

avg-destination-translation-requests

The average number of FW NAT destination translation requests for the selected filter (entity).

avg-destination-translation-request-failures

The average number of FW NAT destination translation request failures for the selected filter (entity).

avg-active-translations

The average number of FW NAT active translations for the selected filter (entity).

avg-backup-pool-translations

The average number of FW NAT backup-pool translations for the selected filter (entity).

avg-log-attempts

The average number of FW NAT log attempts for the selected filter (entity). Only applies to FW NATs in logging mode.

avg-log-failures

The average number of FW NAT log failures for the selected filter (entity). Only applies to FW NATs in logging mode.

avg-active-port-blocks

The average number of FW NAT active port-blocks for the selected filter (entity). Only applies to FW NATs in PBA mode.

avg-port-block-allocations

The average number of FW NAT port-block allocations for the selected filter (entity). Only applies to FW NATs in PBA mode.

avg-port-block-deallocations

The average number of FW NAT port-block de-allocations for the selected filter (entity). Only applies to FW NATs in PBA mode.

avg-zombie-port-blocks-created

The average number of FW NAT zombie port-blocks which have been created for the selected filter (entity). Only applies to FW NATs in PBA mode.

avg-zombie-port-blocks-deleted

The average number of FW NAT zombie port-blocks which have been deleted for the selected filter (entity). Only applies to FW NATs in PBA mode.

avg-active-clients-reached-limit

The average number of FW NAT active-clients limit reached for the selected filter (entity). Only applies to FW NATs in PBA mode.

avg-clients-reached-limit

The average number of FW NAT total-clients limit reached for the selected filter (entity). Only applies to FW NATs in PBA mode.

avg-pcp-request

The average number of FW NAT PCP-requests for the selected filter (entity). Only applies to FW NATs in PCP mode.

avg-pcp-response

The average number of FW NAT PCP-responses for the selected filter (entity). Only applies to FW NATs in PCP mode.

avg-pcp-error

The average number of FW NAT PCP-errors for the selected filter (entity). Only applies to FW NATs in PCP mode.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics fw-nat report(1)

analytics fw-nat scheduled-report

NAME

scheduled-report - Configure scheduled reports for firewall NAT.

MODULE

analytics fw-nat

SYNTAX

Configure the scheduled-report component within the analytics fw-nat module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |
replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[name] | [glob] | [regex]] ...]

DELETE
delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the firewall NAT module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com } frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit 5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the firewall NAT scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics fw-nat report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics fw-nat report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics fw-nat report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics fw-nat scheduled-report(1)

analytics global-settings

NAME

global-settings - set up global analytics parameters.

MODULE

analytics global-settings

SYNTAX

List or set up parameters related to analytics data publishing using the syntax shown in the following sections.

LIST GLOBAL SETTINGS

list analytics global-settings [all-properties | non-default-properties | one-line | recursive]

GLOBAL PUBLISHING SETTINGS

modify avrd-interval [value]

modify avrd-debug-mode [disabled | enabled]

modify disable-all-internal-logging [disabled | enabled]

HSL PUBLISHER

modify external-logging-publisher [publisher name | none]

OFFBOX SETTINGS

modify offbox-protocol [ecm-tm | hsl | none | tcp]

modify offbox-tcp-addresses

[add | delete | none | replace-all-with] { list of IP addresses... }

modify offbox-tcp-port [value]

modify use-offbox [disabled | enabled]

DESCRIPTION

Use the analytics global-settings command to list or modify the following global analytics settings:

• **avrd-interval** - analytics data collection interval in seconds. If this interval is different from the default value (300 seconds), internal statistics are not collected unless avrd-debug-mode is set to enabled. Minimal interval is 20 seconds, maximum interval is 300 seconds.

• **avrd-debug-mode** - enable or disable debug mode. If debug mode is disabled (by default), internal statistics are collected only if avrd-interval is set to the default value (300 seconds).

• **disable-all-internal-logging** - when it is enabled, internal statistics are never collected, regardless of avrd-interval and avrd-debug-mode settings.

• **external-logging-publisher** - choose which logging publisher will be used for the external HSL protocol. Relevant when offbox-protocol is set to 'hsl'.

• **offbox-protocol** - protocol for communication with offbox analytics application. The protocol can be defined as one of the following four options: 'hsl' - (High Speed Logging) a protocol based on UDP or TCP; 'tcp' - the proprietary TCP-based protocol; 'ecm-tm' - a protocol based on Google RPC (grpc); and 'none'. Note: 'ecm-tm' is the recommended option. Note: If 'none' is chosen, the analytics module does not try to connect to the offbox analytics application.

• **offbox-tcp-addresses** - server IP addresses used only if the 'tcp' protocol is chosen. Multiple IP addresses are supported. If more than one IP address is configured, the analytics module first tries to connect to the first IP address in the list. If the connection can't be established, it tries the next IP address, and so on.

• **offbox-tcp-port** - server TCP port for the server IP addresses used only if 'tcp' protocol is chosen. If multiple tcp addresses are configured they all use the same port.

• **use-offbox** - enables and disables all communication with the offbox application on global level.

EXAMPLES

Display all current analytics global settings:

```
list analytics global-settings all-properties
```

Enable the analytics module to communicate with an offbox application using the 'tcp' protocol. In this example, the offbox application IP addresses are 192.168.1.1 and 192.168.1.2, TCP port is - 3344 :

```
modify analytics global-settings offbox-tcp-addresses
```

```
add { 192.168.1.1 192.168.1.2 }
```

```
modify analytics global-settings offbox-tcp-port 3344
```

```
modify analytics global-settings offbox-protocol tcp
modify analytics global-settings use-offbox enabled
```

Create an analytics report every 20 seconds. Note: Since it has a big performance impact, this should be configured for debugging purposes only.

```
modify analytics global-setting avrd-debug-mode enabled
modify analytics global-setting avrd-interval 20
```

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2017. All rights reserved.

BIG-IP 2017-05-14 analytics global-settings(1)

analytics gtm-wideip report

NAME

report - Displays a GTM Wideip analytics report.

MODULE

analytics gtm-wideip

SYNTAX

Show, save or send a analytics gtm-wideip report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ name ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ name ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
field-fmt
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SAVE

```
save report view-by [ name ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ name ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
file [ file name ]
```

```
format [ csv-aggregated | csv-time-series | pdf ]
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
measure [ measure name ]
```

```

    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ name ]
options:
  drilldown {
    {
  entity [ name ]
  values
  {
    [value ...]
  }
  } ...
  }
  email-addresses {
    [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc / desc ]
    } ...
  }
  range [date range]
  smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate GTM Wideip analytics reports. You can generate a GTM Wideip analytics report for the following entities:

• name - GTM wideip name.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics gtm-wideip report view-by name
```

```
show analytics gtm-wideip report view-by name drilldown { { entity name values { some_name } } }
```

```
send-mail analytics gtm-wideip report view-by name measures { dropped } limit 20 order-by { { measure dropped sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

alternate
GTM wideip alternate.

cname-resolutions
GTM wideip cname.

dropped
GTM wideip dropped.

fallback
GTM wideip fallback.

preferred
GTM wideip preferred.

requests
GTM wideip request.

resolutions
GTM wideip resolutions.

return-from-dns
GTM wideip return from DNS.

return-to-dns
GTM wideip return to DNS.

order-by
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-05 analytics gtm-wideip report(1)

analytics http report

NAME

report - Displays an HTTP/L7-DoS analytics report.

MODULE

analytics http

SYNTAX

Show, save or send an analytics http report using the syntax shown in the following sections.

DISPLAY

show report view-by [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet | country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation | os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]

options:

drilldown {
 {

entity [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet | country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation | os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]

values

{
 [value ...]
}

```

} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
}
range [date range]

SAVE
save report view-by [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
  country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
  os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
options:
  drilldown {
    {
entity [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
  country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
  os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
values
{
  [value ...]
}
} ...
}
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
}
range [date range]

SEND
send-mail report view-by [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
  country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
  os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
options:
  drilldown {
    {
entity [ activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
  country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
  os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual ]
values
{
  [value ...]
}
} ...
}
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate HTTP analytics reports. You can generate an HTTP analytics report for the following entities:

- Â· activity-type - Activity type.
- Â· application - Application services.
- Â· attack-id - Application/L7 DoS Attack ID.
- Â· behavioral-signature - Behavioral signature.
- Â· browser - Browser.
- Â· client-ip - A single client identified by an IP address.
- Â· client-subnet - Client subnet.
- Â· country - A country from which HTTP/HTTPS traffic was sent to each of the virtual servers.
- Â· country-code - Country code from which HTTP/HTTPS traffic was sent to each of the virtual servers.
- Â· device-id - Device ID.
- Â· dos-profile - DoS Profile.
- Â· http-method - Method.
- Â· http-transaction-outcome - HTTP Transaction outcomes (Blocked/Dropped/Passthrough/etc.)
- Â· mitigation - Mitigation.
- Â· os - OS name.
- Â· pool-member - Pool members.
- Â· response-code - An HTTP response code that was sent back to the client.
- Â· suspected-ip - Suspected address IP.
- Â· trigger - Trigger.
- Â· url - A URL accessed by HTTP or HTTPS.
- Â· user-agent - A browser identifier sent by the client's browser as part of the request for URL.
- Â· vector - Attack vector.
- Â· virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics http report view-by virtual measures {average-tps} limit 20
```

Gets the average tps of 20 virtual servers (unordered).

```
show analytics http report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps
sort-type desc } }
```

Gets the average tps of the top 20 virtual servers.

```
show analytics http report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps
sort-type desc } } range now-3d--now
```

Gets the average tps of the top 20 virtual servers from the last three days.

```
show analytics http report view-by virtual drilldown { { entity application values { app } } { entity pool-
member values { p1 p2 } } } range now-4d--now-2d measures {average-tps} limit 10 order-by { { measure average-
tps sort-type DESC } }
```

Gets the average tps of the top 10 virtual servers (ordered by average tps) on app iApp (out of several monitored) on pool members p1 and p2 (out of five monitored p1-p5) in the interval ranging from two to four days ago.

```
show analytics http report view-by response-code drilldown { { entity virtual values { v1 } } } measures {
transactions }
```

Gets a distribution of requests per response code on virtual v1.

```
show analytics http report view-by country drilldown { { entity application values { app } } } measures {
average-concurrent-sessions average-sessions } order-by { { measure average-sessions sort-type DESC } } limit
5
```

Gets the new sessions and average concurrent sessions of the top five countries, ordered by the average concurrent sessions on the application app.

```
show analytics http report view-by client-ip drilldown { { entity virtual values { v1 } } } measures { max-
page-load-time } limit 1
```

Gets the client IP address with the worst page load time.

```
show analytics http report view-by application drilldown { { entity pool-member values { p1 p2 } } } measures { transactions } order-by { { measure transactions } } range now-7d--now
```

Gets the distribution of requests per application on pool members p1 and p2 ordered by the number of requests during the last week.

```
save analytics http report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format pdf file report.pdf
```

Gets the average tps of the top 20 virtual servers and exports to a PDF file on the BIG-IP system.

```
save analytics http report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format csv-aggregated file report.csv
```

Gets the average tps of the top 20 virtual servers and exports to a CSV file on the BIG-IP system.

```
save analytics http report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format csv-time-series file report.csv
```

Gets the average tps over time of the top 10 virtual servers and exports to a CSV file on the BIG-IP system.

```
send-mail analytics http report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

Gets the average tps over time of the top 10 virtual servers and sends out an email containing the report as a PDF.

OPTIONS

`device`

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

`device-list`

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

`drilldown`

Specifies specific entities that are used as a filter.

`email-addresses`

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

`file` Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

`format`

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

`include-others`

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

`include-total`

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

`limit`

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

`measures`

Specifies a list of measures that can be used with the chosen entity type. The default value is transactions. The options are:

`average-concurrent-sessions`

The average number of concurrent sessions for each entity.

`average-new-sessions`

The average number of new sessions for each entity.

`average-page-load-time`

The average client page load time for each entity.

`average-request-throughput`

The average request throughput for each entity.

`average-response-throughput`

The average response throughput for each entity.

`average-server-latency`

The average server latency for each entity.

`average-tps`

The average number of transactions per second for each entity.

`client-side-sampled-transactions`

The number of transactions sampled for client side page load time.

max-page-load-time

The maximum client page load time for each entity.

max-request-throughput

The maximum request throughput for each entity.

max-response-throughput

The maximum response throughput for each entity.

max-server-latency

The maximum server latency for each entity.

max-tps

The maximum number of transactions per second for each entity.

transactions

The absolute number of transactions for each entity.

min-server-latency

The minimum server latency for each entity.

average-request-size

The average request size for each entity.

average-response-size

The average response size for each entity.

average-application-response-time

The average application response time for each entity.

min-application-response-time

The minimum application response time for each entity.

max-application-response-time

The maximum application response time for each entity.

average-client-ttfb

The average client TTFB for each entity.

min-client-ttfb

The minimum client TTFB for each entity.

max-client-ttfb

The maximum client TTFB for each entity.

average-clientside-network-latency

The average client-side network latency for each entity.

min-clientside-network-latency

The minimum client-side network latency for each entity.

max-clientside-network-latency

The maximum client-side network latency for each entity.

average-serverside-network-latency

The average server-side network latency for each entity.

min-serverside-network-latency

The minimum server-side network latency for each entity.

max-serverside-network-latency

The maximum server-side network latency for each entity.

average-request-duration

The average request duration for each entity.

min-request-duration

The minimum request duration for each entity.

max-request-duration

The maximum request duration for each entity.

average-response-duration

The average response duration for each entity.

min-response-duration

The minimum response duration for each entity.

max-response-duration

The maximum response duration for each entity.

attacks-count

The total number of attack for each entity.

valid

The total number of valid transactions for each entity.

average-valid-tps

The average number of valid transactions for each entity.

mitigated

The total number of mitigated transaction for each entity.

average-mitigated-tps

The average number of mitigated transaction for each entity.

blocked

The total number of blocked transactions for each entity.

average-blocked-tps

The average number of blocked transactions for each entity.

incomplete

The total number of incomplete transactions for each entity.

average-incomplete-tps

The average number of incomplete transactions for each entity.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The default value for measures is previously chosen measures. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2018-07-05 analytics http report(1)

analytics http scheduled-report

NAME

scheduled-report - Configure scheduled reports for HTTP.

MODULE

analytics http

SYNTAX

Configure the scheduled-report component within the analytics http module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the HTTP module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the HTTP scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics http report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics http report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics http report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics http scheduled-report(1)

analytics ip-intelligence report

NAME

report - Displays an IP Intelligence analytics report.

MODULE

analytics ip-intelligence

SYNTAX

Show, save or send an analytics ip-intelligence report using the syntax shown in the following sections.

DISPLAY

show report view-by [actions | class-name | context-name | context-type | hit-type | policy | protocol | source-ip | vlan]

options:

drilldown {

{

entity [actions | class-name | context-name | context-type | hit-type | policy | protocol | source-ip | vlan]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [actions | class-name | context-name | context-type | hit-type | policy | protocol | source-ip | vlan]

options:

drilldown {

{

entity [actions | class-name | context-name | context-type | hit-type | policy | protocol | source-ip | vlan]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

```

send-mail report view-by [ actions | class-name | context-name | context-type | hit-type | policy | protocol | source-ip | vlan ]
options:
  drilldown {
  {
entity [ actions | class-name | context-name | context-type | hit-type | policy | protocol | source-ip | vlan ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate IP Intelligence analytics reports. You can generate an IP Intelligence analytics report for the following entities:

- Â· actions - Actions (dropped / allowed).
- Â· class-name - Class-Name for the hit.
- Â· context-name - Context Name (VS name, RouteDomain name etc...).
- Â· context-type - Context Type (Virtual, RouteDomain or Global).
- Â· hit-type - Hit Type (BlackList Hit or BlackList+WhiteList Hit).
- Â· policy - Policy Name (Name of the DWBL policy which was hit).
- Â· protocol - Standard IP protocol (tcp, udp, etc...).
- Â· source-ip - Source IPs of hits.
- Â· vlan - Vlan.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics ip-intelligence report view-by actions
```

```
show analytics ip-intelligence report view-by actions drilldown { { entity vlan values { vlan_1 vlan_2 } } }
```

```
send-mail analytics ip-intelligence report view-by actions measures { request-count } limit 20 order-by { {
measure request-count sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

count-class

The total number of events for Dynamic White/Black List with respect to the Class-Names.

request-count

The total number of events for Dynamic White/Black List.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, Itm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-05 analytics ip-intelligence report(1)

analytics ip-intelligence scheduled-report

NAME

scheduled-report - Configure scheduled reports for IP intelligence (DWBL).

MODULE

analytics ip-intelligence

SYNTAX

Configure the scheduled-report component within the analytics ip-intelligence module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE

delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the IP intelligence (DWBL) module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com } frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit 5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the IP intelligence (DWBL) scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics ip-intelligence report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics ip-intelligence report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics ip-intelligence report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics ip-intelligence scheduled-report(1)

analytics ip-layer report

NAME

report - Displays an ip-layer analytics report.

MODULE

analytics ip-layer

SYNTAX

Show, save or send an analytics ip-layer report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ ip ]
options:
  drilldown {
  {
  entity [ ip ]
  values
  {
  [value ...]
  } ...
  }
  field-fmt
  include-total
  include-others
  limit [number of rows]
  measures {
  [measure name ...]
  }
  order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
  }
  range [date range]
```

SAVE

```
save report view-by [ ip ]
options:
  drilldown {
  {
  entity [ ip ]
  values
  {
  [value ...]
  } ...
  }
  file [ file name ]
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
  [measure name ...]
  }
  order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
  }
  range [date range]
```

SEND

```
send-mail report view-by [ ip ]
options:
  drilldown {
  {
```

```

entity [ ip ]
values
{
  [value ...]
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate ip-layer analytics reports. You can generate a IP-layer analytics report for the following entities:

- ip - IP version (IPV4 / IPV6)

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics ip-layer report view-by ip
```

```
show analytics ip-layer report view-by ip drilldown { { entity ip values { IP-V4 } } }
```

```
send-mail analytics ip-layer report view-by ip measures { dropped-pkts } limit 20 order-by { { measure
dropped-pkts sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

transmitted-pkts

The total number of transmitted packets for the selected filter (entity).

received-pkts

The total number of received packets for the selected filter (entity).

dropped-pkts

The total number of dropped packets for the selected filter (entity).

`err-invalid-len`
The total number of error invalid length for the selected filter (entity).

`err-memory`
The total number of error memory for the selected filter (entity).

`err-retransmitted`
The total number of error retransmitted for the selected filter (entity).

`err-protocol`
The total number of error protocol for the selected filter (entity).

`err-options`
The total number of error options for the selected filter (entity).

`err-checksum`
The total number of error checksum for the selected filter (entity).

`received-frags`
The total number of received fragments for the selected filter (entity).

`received_dropped_frags`
The total number of received dropped fragments for the selected filter (entity).

`transmitted-frags`
The total number of transmitted fragments for the selected filter (entity).

`transmitted-dropped-frags`
The total number of transmitted dropped fragment for the selected filter (entity).

`reassembled-frags`
The total number of reassembled fragment for the selected filter (entity).

`reassembled-dropped-frags`
The total number of 'too long' reassembled dropped fragment for the selected filter (entity).

`order-by`
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

`range`
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

`smtp-config-override`
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

`show`, `save`, `send-mail`, `tms`, `itm profile analytics`, `analytics`, `analytics report`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-07-22 analytics ip-layer report(1)

analytics ip-layer scheduled-report

NAME

scheduled-report - Configure scheduled reports for IP layer.

MODULE

analytics ip-layer

SYNTAX

Configure the scheduled-report component within the analytics ip-layer module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

```
replace-all-with { email-address [string] }
first-time [date]
frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
include-total [enabled | disabled]
multi-leveled-report {
chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
limit [number of rows]
time-diff [last-hour | last-day | last-week | last-month | last-year]
view-by { entity name [string] }
measures [none | add | delete | modify | replace-all-with] { measure name [string] }
}
predefined-report-name [name]
smtp-config [name]
device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the IP layer module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the IP layer scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics ip-layer report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics ip-layer report.

measures

The measures which are available for the selected entities.

`predefined-report-name`
Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with `multi-leveled-report`.

`smtp-config`
Defines which SMTP configuration will be used to send the scheduled report. If set to `none`, the scheduled report will be disabled.

`device-group`
Defines the device-group which the report should generate the report for. If `'none'` is set to this field, then the report will be generate for the `'self'` device.

SEE ALSO

`list`, `modify`, `show`, `tms`, `analytics ip-layer report`, `sys smtp-server`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics ip-layer scheduled-report(1)

analytics lsn-pool report

NAME

`report` - Displays an LSN Pool analytics report.

MODULE

`analytics lsn-pool`

SYNTAX

Show, save or send an analytics lsn-pool report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ lsn-pool ]
```

```
options:
```

```
  drilldown {
```

```
  {
```

```
    entity [ none ]
```

```
  values
```

```
  {
```

```
    [value ...]
```

```
  }
```

```
  } ...
```

```
}
```

```
field-fmt
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
  [measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
  measure [ measure name ]
```

```
  sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SAVE

```
save report view-by [ lsn-pool ]
```

```
options:
```

```
  drilldown {
```

```
  {
```

```
    entity [ none ]
```

```
  values
```

```
  {
```

```
    [value ...]
```

```
  }
```

```
  } ...
```

```
}
```

```
file [ file name ]
```

```
format [ csv-aggregated | csv-time-series | pdf ]
```

```
include-total
```

```

include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ lsn-pool ]
options:
drilldown {
  {
entity [ none ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate lsn-pool analytics reports. You can generate a LSN analytics report for the following entities:

- lsn-pool - LSN Pool translation statistics

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics lsn-pool report view-by lsn-pool
```

```
send-mail analytics lsn-pool report view-by lsn-pool measures { active-translations } limit 20 order-by { {
measure active-translations sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

active-translations

The average number of active LSN translations for the selected filter (entity).

translation-request

The total number of LSN translation requests for the selected filter (entity).

translation-failure

The total number of failed LSN translations for the selected filter (entity).

translation-from-backup

The total number of LSN translations using the backup pool for selected filter (entity). Only applies to LSN Pool in Deterministic NAT mode.

active-pb

The average number of active port-blocks for the selected filter (entity). Only applies to LSN Pool in PBA mode.

pb-allocations

The total number of port-block allocations for the selected filter (entity). Only applies to LSN Pool in PBA mode.

pb-freed

The total number of port-block deallocations for the selected filter (entity). Only applies to LSN Pool in PBA mode.

zombie-block-created

The total number of zombie blocks created for the selected filter (entity). Only applies to LSN Pool in PBA mode.

zombie-block-deleted

The total number of zombie blocks deleted for the selected filter (entity). Only applies to LSN Pool in PBA mode.

active-pb-clients-reached-limit

The average number of port-block clients that have reached the port-block limit for the selected filter (entity). Only applies to LSN Pool in PBA mode.

pb-client-reached-limit

The total number of port-block clients that have reached the port-block limit for the selected filter (entity). Only applies to LSN Pool in PBA mode.

pcp-requests

The total number of PCP requests for the selected filter (entity). Only applies to LSN Pool with PCP profile.

pcp-responses

The total number of PCP responses for the selected filter (entity). Only applies to LSN Pool with PCP profile.

pcp-errors

The total number of PCP errors for the selected filter (entity). Only applies to LSN Pool with PCP profile.

log-attempts

The total number of logging attempts for the selected filter (entity). Only applies to LSN Pool with Log publisher attached.

log-failures

The total number of logging failures for the selected filter (entity). Only applies to LSN Pool with Log publisher attached.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 analytics lsn-pool report(1)

analytics lsn-pool scheduled-report

NAME

scheduled-report - Configure scheduled reports for LSN pool.

MODULE

analytics lsn-pool

SYNTAX

Configure the scheduled-report component within the analytics lsn-pool module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the LSN pool module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the LSN pool scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics lsn-pool report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics lsn-pool report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics lsn-pool report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics lsn-pool scheduled-report(1)

analytics memory-per-process report

NAME

report - Displays a Memory Per Process analytics report.

MODULE

analytics memory-per-process

SYNTAX

Show, save or send a analytics memory-per-process report using the syntax shown in the following sections.

DISPLAY

show report view-by [pid | proc-name | slot]

options:

drilldown {

{

entity [pid | proc-name | slot]

values

{

[value ...]

}

```

} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SAVE
save report view-by [ pid | proc-name | slot ]
options:
  drilldown {
    {
entity [ pid | proc-name | slot ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ pid | proc-name | slot ]
options:
  drilldown {
    {
entity [ pid | proc-name | slot ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate Memory Per Process analytics reports. You can generate a Memory Per Process analytics report for the following entities:

• pid - Process ID.

• proc-name - Process name.

• slot - Slot ID.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics memory-per-process report view-by pid
```

```
show analytics memory-per-process report view-by pid drilldown { { entity slot values { 0 1 } } }
```

```
send-mail analytics memory-per-process report view-by pid measures { vsz } limit 20 order-by { { measure vsz  
sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

rss Resident set size.

vsz Virtual memory size.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics memory-per-process report(1)

analytics memory report

NAME

report - Displays an memory analytics report.

MODULE
analytics memory

SYNTAX
Show, save or send an analytics memory report using the syntax shown in the following sections.

DISPLAY
show report view-by [slot]
options:
drilldown {
 {
entity [slot]
values
{
 [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
 [measure name ...]
}
order-by {
 {
 measure [measure name]
 sort-type [asc / desc]
 } ...
}
range [date range]

SAVE
save report view-by [slot]
options:
drilldown {
 {
entity [slot]
values
{
 [value ...]
}
} ...
}
file [file name]
format [csv-aggregated | csv-time-series | pdf]
include-total
include-others
limit [number of rows]
measures {
 [measure name ...]
}
order-by {
 {
 measure [measure name]
 sort-type [asc / desc]
 } ...
}
range [date range]

SEND
send-mail report view-by [slot]
options:
drilldown {
 {
entity [slot]
values
{
 [value ...]
}
} ...
}
email-addresses {
 [email address ...]
}
format [csv-aggregated | csv-time-series | pdf]
include-total
include-others
limit [number of rows]
measures {
 [measure name ...]
}
order-by {
 {
 measure [measure name]

```
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate memory analytics reports. You can generate a memory analytics report for the following entities:

• slot - Slot ID

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics memory report view-by slot
```

```
show analytics memory report view-by slot drilldown { { entity slot values { IP-V4 } } }
```

```
send-mail analytics memory report view-by slot measures { dropped-pkts } limit 20 order-by { { measure host-usage sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

tmm-total-ram The average of total TMM memory for the selected filter (entity).

tmm-used-ram The average of used TMM memory for the selected filter (entity).

tmm-free-ram The average of free TMM memory for the selected filter (entity).

other-total-ram The average of total other (daemons and kernel) memory for the selected filter (entity).

other-used-ram The average of used other (daemons and kernel) memory for the selected filter (entity).

other-free-ram The average of free other (daemons and kernel) memory for the selected filter (entity).

system-total-ram The average of total system memory for the selected filter (entity).

system-used-ram The average of used system memory for the selected filter (entity).

system-free-ram The average of free system memory for the selected filter (entity).

swap-total-ram The average of total swap memory for the selected filter (entity).

swap-used-ram The average of used swap memory for the selected filter (entity).

swap-free-ram The average of free swap memory for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-07-22 analytics memory report(1)

analytics memory scheduled-report

NAME

scheduled-report - Configure scheduled reports for memory.

MODULE

analytics memory

SYNTAX

Configure the scheduled-report component within the analytics memory module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the memory module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the memory scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics memory report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics memory report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics memory report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics memory scheduled-report(1)

analytics network report

NAME

report - Displays a network firewall analytics report.

MODULE

analytics network

SYNTAX

Show, save or send an analytics network report using the syntax shown in the following sections.

DISPLAY

show report view-by [I314-errors-error-reason | I314-errors-network-protocol | I314-errors-action | I314-errors-source-ip | I314-errors-destination-
acl-enforced-application | acl-enforced-destination-ip | acl-enforced-destination-port | acl-enforced-policy | acl-enforced-rule |
acl-enforced-rule-action | acl-enforced-rule-context | acl-enforced-rule-context-type | acl-enforced-self-ip | acl-enforced-server-ip |
acl-enforced-source-ip | acl-enforced-source-port | acl-enforced-translation-pool | acl-enforced-translation-type | acl-enforced-vlan |
acl-mgmt-application | acl-mgmt-destination-ip | acl-mgmt-destination-port | acl-mgmt-rule | acl-mgmt-rule-action |
acl-mgmt-rule-context | acl-mgmt-source-ip | acl-mgmt-source-port |
acl-staged-application | acl-staged-destination-ip | acl-staged-destination-port | acl-staged-policy | acl-staged-rule |

```

    acl-staged-rule-action | acl-staged-rule-context | acl-staged-rule-context-type | acl-staged-self-ip | acl-staged-server-ip |
    acl-staged-source-ip | acl-staged-source-port | acl-staged-translation-pool | acl-staged-translation-type | acl-staged-vlan ]
options:
  drilldown {
  {
entity [ I314-errors-error-reason | I314-errors-network-protocol | I314-errors-action | I314-errors-source-ip | I314-errors-destination-ip | I314-errors-
acl-enforced-application | acl-enforced-destination-ip | acl-enforced-destination-port | acl-enforced-policy | acl-enforced-rule |
acl-enforced-rule-action | acl-enforced-rule-context | acl-enforced-rule-context-type | acl-enforced-self-ip | acl-enforced-server-ip |
acl-enforced-source-ip | acl-enforced-source-port | acl-enforced-translation-pool | acl-enforced-translation-type | acl-enforced-vlan |
acl-mgmt-application | acl-mgmt-destination-ip | acl-mgmt-destination-port | acl-mgmt-rule | acl-mgmt-rule-action |
acl-mgmt-rule-context | acl-mgmt-source-ip | acl-mgmt-source-port |
acl-staged-application | acl-staged-destination-ip | acl-staged-destination-port | acl-staged-policy | acl-staged-rule |
acl-staged-rule-action | acl-staged-rule-context | acl-staged-rule-context-type | acl-staged-self-ip | acl-staged-server-ip |
acl-staged-source-ip | acl-staged-source-port | acl-staged-translation-pool | acl-staged-translation-type | acl-staged-vlan ]
values
{
[value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc | desc ]
} ...
}
}
range [date range]

SAVE
save report view-by [ I314-errors-error-reason | I314-errors-network-protocol | I314-errors-action | I314-errors-source-ip | I314-errors-destination-i
acl-enforced-application | acl-enforced-destination-ip | acl-enforced-destination-port | acl-enforced-policy | acl-enforced-rule |
acl-enforced-rule-action | acl-enforced-rule-context | acl-enforced-rule-context-type | acl-enforced-self-ip | acl-enforced-server-ip |
acl-enforced-source-ip | acl-enforced-source-port | acl-enforced-translation-pool | acl-enforced-translation-type | acl-enforced-vlan |
acl-mgmt-application | acl-mgmt-destination-ip | acl-mgmt-destination-port | acl-mgmt-rule | acl-mgmt-rule-action |
acl-mgmt-rule-context | acl-mgmt-source-ip | acl-mgmt-source-port |
acl-staged-application | acl-staged-destination-ip | acl-staged-destination-port | acl-staged-policy | acl-staged-rule |
acl-staged-rule-action | acl-staged-rule-context | acl-staged-rule-context-type | acl-staged-self-ip | acl-staged-server-ip |
acl-staged-source-ip | acl-staged-source-port | acl-staged-translation-pool | acl-staged-translation-type | acl-staged-vlan ]
options:
  drilldown {
  {
entity [ I314-errors-error-reason | I314-errors-network-protocol | I314-errors-action | I314-errors-source-ip | I314-errors-destination-ip | I314-errors-
acl-enforced-application | acl-enforced-destination-ip | acl-enforced-destination-port | acl-enforced-policy | acl-enforced-rule |
acl-enforced-rule-action | acl-enforced-rule-context | acl-enforced-rule-context-type | acl-enforced-self-ip | acl-enforced-server-ip |
acl-enforced-source-ip | acl-enforced-source-port | acl-enforced-translation-pool | acl-enforced-translation-type | acl-enforced-vlan |
acl-mgmt-application | acl-mgmt-destination-ip | acl-mgmt-destination-port | acl-mgmt-rule | acl-mgmt-rule-action |
acl-mgmt-rule-context | acl-mgmt-source-ip | acl-mgmt-source-port |
acl-staged-application | acl-staged-destination-ip | acl-staged-destination-port | acl-staged-policy | acl-staged-rule |
acl-staged-rule-action | acl-staged-rule-context | acl-staged-rule-context-type | acl-staged-self-ip | acl-staged-server-ip |
acl-staged-source-ip | acl-staged-source-port | acl-staged-translation-pool | acl-staged-translation-type | acl-staged-vlan ]
values
{
[value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc | desc ]
} ...
}
}
range [date range]

SEND
send-mail report view-by [ I314-errors-error-reason | I314-errors-network-protocol | I314-errors-action | I314-errors-source-ip | I314-errors-destina
acl-enforced-application | acl-enforced-destination-ip | acl-enforced-destination-port | acl-enforced-policy | acl-enforced-rule |
acl-enforced-rule-action | acl-enforced-rule-context | acl-enforced-rule-context-type | acl-enforced-self-ip | acl-enforced-server-ip |
acl-enforced-source-ip | acl-enforced-source-port | acl-enforced-translation-pool | acl-enforced-translation-type | acl-enforced-vlan |
acl-mgmt-application | acl-mgmt-destination-ip | acl-mgmt-destination-port | acl-mgmt-rule | acl-mgmt-rule-action |
acl-mgmt-rule-context | acl-mgmt-source-ip | acl-mgmt-source-port |
acl-staged-application | acl-staged-destination-ip | acl-staged-destination-port | acl-staged-policy | acl-staged-rule |

```

```

acl-staged-rule-action | acl-staged-rule-context | acl-staged-rule-context-type | acl-staged-self-ip | acl-staged-server-ip |
acl-staged-source-ip | acl-staged-source-port | acl-staged-translation-pool | acl-staged-translation-type | acl-staged-vlan ]
options:
drilldown {
{
entity [ I314-errors-error-reason | I314-errors-network-protocol | I314-errors-action | I314-errors-source-ip | I314-errors-destination-ip | I314-errors-
acl-enforced-application | acl-enforced-destination-ip | acl-enforced-destination-port | acl-enforced-policy | acl-enforced-rule |
acl-enforced-rule-action | acl-enforced-rule-context | acl-enforced-rule-context-type | acl-enforced-self-ip | acl-enforced-server-ip |
acl-enforced-source-ip | acl-enforced-source-port | acl-enforced-translation-pool | acl-enforced-translation-type | acl-enforced-vlan |
acl-mgmt-application | acl-mgmt-destination-ip | acl-mgmt-destination-port | acl-mgmt-rule | acl-mgmt-rule-action |
acl-mgmt-rule-context | acl-mgmt-source-ip | acl-mgmt-source-port |
acl-staged-application | acl-staged-destination-ip | acl-staged-destination-port | acl-staged-policy | acl-staged-rule |
acl-staged-rule-action | acl-staged-rule-context | acl-staged-rule-context-type | acl-staged-self-ip | acl-staged-server-ip |
acl-staged-source-ip | acl-staged-source-port | acl-staged-translation-pool | acl-staged-translation-type | acl-staged-vlan ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc | desc ]
} ...
}
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate network firewall analytics reports. You can generate a network firewall analytics report for the following entities:

- Â· action - Action taken (allowed/dropped).
- Â· acl-enforced-application - Application services (ACL - Enforced).
- Â· acl-enforced-destination-ip - Destination IP Address (ACL - Enforced).
- Â· acl-enforced-destination-port - Destination IP Port (ACL - Enforced).
- Â· acl-enforced-policy - Policy (ACL - Enforced).
- Â· acl-enforced-rule-action - Rule Action (ACL - Enforced).
- Â· acl-enforced-rule-context - Rule Context (ACL - Enforced).
- Â· acl-enforced-rule-context-type - Rule Context Type (ACL - Enforced).
- Â· acl-enforced-rule - Rule (ACL - Enforced).
- Â· acl-enforced-self-ip - Self IP Address (ACL - Enforced).
- Â· acl-enforced-server-ip - Server IP Address (ACL - Enforced).
- Â· acl-enforced-source-ip - Source IP Address (ACL - Enforced).
- Â· acl-enforced-source-port - Source IP Port (ACL - Enforced).
- Â· acl-enforced-translation-pool - Translation Pool (ACL - Enforced).
- Â· acl-enforced-translation-type - Translation Type (ACL - Enforced).
- Â· acl-enforced-vlan - VLAN (ACL - Enforced).
- Â· acl-mgmt-application - Application services (ACL - Management).
- Â· acl-mgmt-destination-ip - Destination IP Address (ACL - Management).
- Â· acl-mgmt-destination-port - Destination IP Port (ACL - Management).
- Â· acl-mgmt-rule-action - Rule Action (ACL - Management).
- Â· acl-mgmt-rule-context - Rule Context (ACL - Management).
- Â· acl-mgmt-rule - Rule (ACL - Management).

- Â· acl-mgmt-source-ip - Source IP Address (ACL - Management).
- Â· acl-mgmt-source-port - Source IP Port (ACL - Management).
- Â· acl-staged-application - Application services (ACL - Staged).
- Â· acl-staged-destination-ip - Destination IP Address (ACL - Staged).
- Â· acl-staged-destination-port - Destination IP Port (ACL - Staged).
- Â· acl-staged-policy - Policy (ACL - Staged).
- Â· acl-staged-rule-action - Rule Action (ACL - Staged).
- Â· acl-staged-rule-context - Rule Context (ACL - Staged).
- Â· acl-staged-rule-context-type - Rule Context Type (ACL - Staged).
- Â· acl-staged-rule - Rule (ACL - Staged).
- Â· acl-staged-self-ip - Self IP Address (ACL - Staged).
- Â· acl-staged-server-ip - Server IP Address (ACL - Staged).
- Â· acl-staged-source-ip - Source IP Address (ACL - Staged).
- Â· acl-staged-source-port - Source IP Port (ACL - Staged).
- Â· acl-staged-translation-pool - Translation Reason (ACL - Staged).
- Â· acl-staged-translation-type - Translation Type (ACL - Staged).
- Â· acl-staged-vlan - VLAN (ACL - Staged).
- Â· I314-errors-action - Network firewall errors action.
- Â· I314-errors-destination-ip - Destination IP address (Network firewall errors).
- Â· I314-errors-error-reason - Network firewall error reason.
- Â· I314-errors-network-protocol - Destination port (Network protocol).
- Â· I314-errors-source-ip - Source IP address (Network firewall errors).
- Â· I314-errors-vlan - VLAN (Network firewall errors).

EXAMPLES

show analytics network report view-by acl-enforced-rule

show analytics network report view-by acl-staged-vlan drilldown { { entity acl-staged-destination-port values { 80 } } }

send-mail analytics network report view-by acl-mgmt-source-ip limit 20 format pdf email-addresses { some.one@someaddress.com }

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not

including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

acl-matches

The total number of ACL rule matches. Applicable only to view-by entities starting with "acl-".

errors

The total number of firewall errors. Applicable only to view-by entities starting with "l3l3-errors-".

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-10 analytics network report(1)

analytics network scheduled-report

NAME

scheduled-report - Configure scheduled reports for network.

MODULE

analytics network

SYNTAX

Configure the scheduled-report component within the analytics network module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE

delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the network module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com } frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit 5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the network scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics network report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics network report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics network report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

analytics network stale-rules

NAME

stale-rules - Displays a network firewall stale rules report.

MODULE

analytics network

SYNTAX

Show an analytics network stale-rules report using the syntax shown in the following sections.

DISPLAY

```
show stale-rules type [ enforced | staged ]
```

options:

```
drilldown {
  {
  entity [ context | policy | rule-name ]
  values
  {
  [value ...]
  }
  } ...
}
```

field-fmt

```
first-rule-number [ value ]
number-of-rules [ value ]
range [ date range ]
```

DESCRIPTION

Use this command to generate network firewall stale rules reports. A stale rule is one that has had not hits, or very few hits, over a specified time period. The report is displayed in order from the least-hit rules (including rules with no hits) to the most hit rules. You can generate a stale rules report for either enforced or staged rules.

EXAMPLES

```
show analytics network stale-rules type enforced
```

Shows a stale rules report for enforced rules (either inline or not).

```
show analytics network stale-rules type staged drilldown { { entity context values { /Common/virtual_server_1 } } }
```

Shows a stale rules report for staged rules in the context of the virtual server /Common/virtual_server_1

```
show analytics network stale-rules type enforced number-of-rules 100 range now-1w
```

Shows a stale rules report for enforced rules. 100 rules are shown in the report. This report is shown for the last week (including the last day).

```
show analytics network stale-rules type enforced first-rule-number 10 number-of-rules 100 range now-1w
```

Shows a stale rules report for enforced rules. The first least hit 9 rules are skipped, and 100 rules are shown in the report. This report is shown for the last week (including the last day).

```
show analytics network stale-rules type enforced first-rule-number 10 number-of-rules 100 range now-1d--now-1w
```

Shows a stale rules report for enforced rules. The first least hit 9 rules are skipped, and 100 rules are shown in the report. This report is shown for the last week, excluding the last day.

OPTIONS

drilldown

Specifies specific entities that are used as a filter.

field-fmt

Shows statistics in field format for the specified items.

first-rule-number

Specifies the first rule number being displayed (rules are ordered by hit count in an ascending order).

number-of-rules

Specifies the maximum number of firewall rules being displayed in the output result set. The default value is 10.

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

SEE ALSO

analytics, analytics report, security analytics settings, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-11-04 analytics network stale-rules(1)

analytics pem report

NAME

report - Displays an pem analytics report.

MODULE

analytics pem

SYNTAX

Show, save or send an analytics pem report using the syntax shown in the following sections.

DISPLAY

show report view-by [application | category | url-category | policy | service | action | tower | subscribers | subscriber-name | ip-list | device-na

options:

```
drilldown {
  {
  entity [ application | category | url-category | policy | service | action | tower | subscribers | subscriber-name | ip-list | device-name | device-o
values
  {
  [value ...]
  }
  } ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SAVE

save report view-by [application | category | url-category | policy | service | action | tower | subscribers | subscriber-name | ip-list | device-na

options:

```
drilldown {
  {
  entity [ application | category | url-category | policy | service | action | tower | subscribers | subscriber-name | ip-list | device-name | device-o
values
  {
  [value ...]
  }
  } ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
}
range [date range]
```

```

SEND
send-mail report view-by [ application | category | url-category | policy | service | action | tower | subscribers| subscriber-name | ip-list | device-os ]
options:
  drilldown {
    {
  entity [ application | category | url-category | policy | service | action | tower | subscribers | subscriber-name | ip-list | device-name | device-os ]
  values
  {
    [value ...]
  } ...
  }
  email-addresses {
    [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc / desc ]
    } ...
  }
  range [date range]
  smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate pem analytics reports. You can generate a PEM analytics report for the following entities:

- Â· application - Classification application
- Â· category - Classification category
- Â· url-category - URL category
- Â· policy - Classification application policy
- Â· service - Forwarding service (endpoint)
- Â· action - PEM action (Gate, Forward, ICAP, Modify HTTP header, etc...)
- Â· tower - Tower the subscriber communicates from
- Â· subscribers - Subscriber summary statistics
- Â· subscriber-name - Subscriber name
- Â· ip-list - IP address list
- Â· device-name - Device name
- Â· device-os - Device operating system
- Â· called-station - Called station
- Â· calling-station - Calling station
- Â· subscriber-type - Subscriber type
- Â· user-name - User Name

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics pem report view-by category
```

```
show analytics pem report view-by category drilldown { { entity policy values { Some_Policy_Name } } }
```

```
send-mail analytics pem report view-by tower measures { total_bytes_in } limit 20 order-by { { measure total_bytes_in sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

hitcount

The total number of classified flows for the selected filter (entity).

total-bytes-in

The total number of bytes received for the selected filter (entity).

total-bytes-out

The total number of bytes sent for the selected filter (entity).

total-flows-opened

The total number of flows (classified and non-classified) for the selected filter (entity).

total-flows-closed

The total number of closed flows for the selected filter (entity).

total-subscribers-login

The total number of login events for the selected filter (entity).

total-subscribers-logout

The total number of logout events for the selected filter (entity).

avg-distinct-apps

The average number of distinct applications for the selected filter (entity).

avg-distinct-categories

The average number of distinct categories for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-07-22 analytics pem report(1)

NAME

scheduled-report - Configure scheduled reports for PEM.

MODULE

analytics pem

SYNTAX

Configure the scheduled-report component within the analytics pem module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |

replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE

delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the PEM module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the PEM scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart

path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics pem report.

limit
The number of view-by entities displayed in the scheduled report.

time-diff
The time range for the report.

view-by
The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics pem report.

measures
The measures which are available for the selected entities.

predefined-report-name
Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config
Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group
Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics pem report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics pem scheduled-report(1)

analytics pool-traffic report

NAME
report - Displays a Pool Traffic Stats analytics report.

MODULE
analytics pool-traffic

SYNTAX
Show, save or send a analytics pool-traffic report using the syntax shown in the following sections.

```
DISPLAY
show report view-by [ pool-address | pool-name ]
options:
drilldown {
{
entity [ pool-address | pool-name ]
values
{
[value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
```

```

SAVE
save report view-by [ pool-address | pool-name ]
options:
  drilldown {
  {
entity [ pool-address | pool-name ]
values
{
[value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]

```

```

SEND
send-mail report view-by [ pool-address | pool-name ]
options:
  drilldown {
  {
entity [ pool-address | pool-name ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate Pool Traffic Stats analytics reports. You can generate a Pool Traffic Stats analytics report for the following entities:

Â· pool-address - Pool Address.

Â· pool-name - Pool Name.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics pool-traffic report view-by pool-address
```

```
show analytics pool-traffic report view-by pool-address drilldown { { entity pool-name values { pool_name } } }
```

```
send-mail analytics pool-traffic report view-by pool-address measures { server-bits-out } limit 20 order-by { { measure server-bits-out sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

max-server-concurrent-connections

Maximum Concurrent connections at peak.

server-bits-in

Incoming bits.

server-bits-out

Outgoing bits.

server-concurrent-connections

Server Concurrent connections.

server-connections

Server Connections.

server-packets-in

Incoming packets.

server-packets-out

Outgoing packets.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics pool-traffic report(1)

analytics pool-traffic scheduled-report

NAME

scheduled-report - Configure scheduled reports for pool traffic.

MODULE

analytics pool-traffic

SYNTAX

Configure the scheduled-report component within the analytics pool-traffic module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the pool traffic module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the pool traffic scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics pool-traffic report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics pool-traffic report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics pool-traffic report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics pool-traffic scheduled-report(1)

analytics proc-cpu report

NAME

report - Displays a Process CPU Utilization analytics report.

MODULE

analytics proc-cpu

SYNTAX

Show, save or send a analytics proc-cpu report using the syntax shown in the following sections.

DISPLAY

show report view-by [blade-num | process-id | process-name]

options:

drilldown {

{

entity [blade-num | process-id | process-name]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [blade-num | process-id | process-name]

options:

```

drilldown {
  {
entity [ blade-num | process-id | process-name ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

```

SEND
send-mail report view-by [ blade-num | process-id | process-name ]
options:
drilldown {
  {
entity [ blade-num | process-id | process-name ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [numbers of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate Process CPU Utilization analytics reports. You can generate a Process CPU Utilization analytics report for the following entities:

• blade-num - Blade Number.

• process-id - Process ID.

• process-name - Process Name.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics proc-cpu report view-by blade-num
```

```
show analytics proc-cpu report view-by process-name drilldown { { entity blade-num values { 0 } } }
```

```
send-mail analytics proc-cpu report view-by blade-num measures { cpu-usage } limit 20 order-by { { measure
cpu-usage sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device
Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list
Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown
Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

cpu-usage

CPU Usage.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics proc-cpu report(1)

analytics proc-cpu scheduled-report

NAME

scheduled-report - Configure scheduled reports for process CPU.

MODULE

analytics proc-cpu

SYNTAX

Configure the scheduled-report component within the analytics proc-cpu module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |
replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

```
view-by { entity name [string] }
measures [none | add | delete | modify | replace-all-with] { measure name [string] }
}
predefined-report-name [name]
smtp-config [name]
device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the process CPU module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the process CPU scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics proc-cpu report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics proc-cpu report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the

scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics proc-cpu report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics proc-cpu scheduled-report(1)

analytics protocol-inspection report

NAME

report - Displays an IPS analytics report.

MODULE

analytics protocol-inspection

SYNTAX

Show, save or send an analytics protocol-inspection report using the syntax shown in the following sections.

DISPLAY

show report view-by [accuracy | action | attack-type | dest-ip | dest-port | inspection-id | inspection-name | perfimpact | profile | protocol | reference | risk | service | source-country | src-ip | subscriber | support-id | virtual | vlan]

options:

drilldown {

{

entity [accuracy | action | attack-type | dest-ip | dest-port | inspection-id | inspection-name | perfimpact | profile | protocol | reference | risk | service | source-country | src-ip | subscriber | support-id | virtual | vlan]

values

{
[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [accuracy | action | attack-type | dest-ip | dest-port | inspection-id | inspection-name | perfimpact | profile | protocol | reference | risk | service | source-country | src-ip | subscriber | support-id | virtual | vlan]

options:

drilldown {

{

entity [accuracy | action | attack-type | dest-ip | dest-port | inspection-id | inspection-name | perfimpact | profile | protocol | reference | risk | service | source-country | src-ip | subscriber | support-id | virtual | vlan]

values

{
[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

```

include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ accuracy | action | attack-type | dest-ip | dest-port | inspection-id | inspection-name | perfimpact |
  profile | protocol | reference | risk | service | source-country | src-ip |
  subscriber | support-id | virtual | vlan ]
options:
drilldown {
  {
entity [ accuracy | action | attack-type | dest-ip | dest-port | inspection-id | inspection-name | perfimpact |
  profile | protocol | reference | risk | service | source-country | src-ip |
  subscriber | support-id | virtual | vlan ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate IPS analytics reports. You can generate an IPS analytics report for the following entities:

- Â· accuracy - Accuracy.
- Â· action - Action.
- Â· attack-type - Attack type of the illegal request.
- Â· dest-ip - Destination IP addresses of requests.
- Â· dest-port - Destination port.
- Â· inspection-id - Inspection Id.
- Â· inspection-name - Inspection name.
- Â· perfimpact - Performance Impact.
- Â· profile - Profile.
- Â· protocol - Protocol.
- Â· reference - Industry Reference.
- Â· risk - Risk.
- Â· service - Service.
- Â· source-country - Country from which the traffic originated.
- Â· src-ip - Source IP addresses of requests.
- Â· subscriber - Subscriber names.
- Â· support-id - Support Id.
- Â· virtual - Virtual servers.

• vlan - Vlan.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics protocol-inspection report view-by accuracy
```

```
show analytics protocol-inspection report view-by accuracy drilldown { { entity virtual values { virtual_1 } } }
```

```
send-mail analytics protocol-inspection report view-by accuracy measures { last-updated } limit 20 order-by { { measure last-updated sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

dest-ip-count

Destination IP count.

last-updated

Last Updated.

occurrences

Occurrence count.

source-ip-count

Source IP count.

virtual-count

Virtual count.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

analytics protocol-security-http report

NAME

report - Displays a HTTP Protocol Security analytics report.

MODULE

analytics protocol-security-http

SYNTAX

Show, save or send a analytics protocol-security-http report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ action | application | client-ip | url | violation | virtual ]
options:
  drilldown {
  {
entity [ action | application | client-ip | url | violation | virtual ]
values
{
  [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SAVE

```
save report view-by [ action | application | client-ip | url | violation | virtual ]
options:
  drilldown {
  {
entity [ action | application | client-ip | url | violation | virtual ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SEND

```
send-mail report view-by [ action | application | client-ip | url | violation | virtual ]
options:
  drilldown {
  {
entity [ action | application | client-ip | url | violation | virtual ]
values
{
  [value ...]
```

```

}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate HTTP Protocol Security analytics reports. You can generate a HTTP Protocol Security analytics report for the following entities:

- Â· action - Enforcement action on request(Legal or alarmed/Blocked/Dropped).
- Â· application - Application services.
- Â· client-ip - A single client identified by an IP address.
- Â· url - A URL accessed by HTTP or HTTPS.
- Â· violation - Violation types.
- Â· virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics protocol-security-http report view-by action
```

```
show analytics protocol-security-http report view-by action drilldown { { entity application values {
application_name } } }
```

```
send-mail analytics protocol-security-http report view-by action measures { requests } limit 20 order-by { {
measure requests sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

occurrences

Occurrence count.

requests
Requests count.

order-by
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics protocol-security-http report(1)

analytics protocol-security-http scheduled-report

NAME

scheduled-report - Configure scheduled reports for protocol security HTTP (PSM HTTP).

MODULE

analytics protocol-security-http

SYNTAX

Configure the scheduled-report component within the analytics protocol-security-http module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the protocol security HTTP (PSM HTTP) module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from being generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the protocol security HTTP (PSM HTTP) scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics protocol-security-http report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics protocol-security-http report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics protocol-security-http report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics protocol-security-http scheduled-report(1)

analytics protocol-security report

NAME

report - Displays a Protocol Security analytics report.

MODULE

analytics protocol-security

SYNTAX

Show, save, or send an analytics protocol-security report using the syntax shown in the following sections.

DISPLAY

show report view-by [application | virtual-server | ip | violation| request-type | protocol-type]

options:

drilldown {

{

entity [application | virtual-server | ip | violation| request-type | protocol-type]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [application | virtual-server | ip | violation| request-type | protocol-type]

options:

drilldown {

{

entity [application | virtual-server | ip | violation| request-type | protocol-type]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [application | virtual-server | ip | violation| request-type | protocol-type]

options:

drilldown {

{

entity [application | virtual-server | ip | violation| request-type | protocol-type]

values

{

[value ...]

}

} ...

}

email-addresses {

[email address ...]

}

format [csv-aggregated | csv-time-series | pdf]

include-total

```

include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate protocol-security analytics reports. You can generate a protocol-security analytics report for the following entities:

- Â· application - Application services.
- Â· virtual-server - Virtual servers.
- Â· ip - Source IP addresses.
- Â· violation - Violation types.
- Â· protocol-type - Protocol type (HTTP/FTP/SMTP)
- Â· request-type - PRequest type (Legal or Alarmed/Blocked/Dropped)

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics protocol-security report view-by protocol-type
```

```
show analytics protocol-security report view-by request-type drilldown { { entity protocol-type values { HTTP } } }
```

```
send-mail analytics protocol-security report view-by protocol-type measures {transactions} limit 20 order-by { { measure transactions sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the manpage for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

requests

Request count.

occurrences

Number of occurrences for the selected filter (relevant for violation entity only)

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is

desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, itm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 analytics protocol-security report(1)

analytics protocol-security scheduled-report

NAME

scheduled-report - Configure scheduled reports for protocol security (PSM).

MODULE

analytics protocol-security

SYNTAX

Configure the scheduled-report component within the analytics protocol-security module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the protocol security (PSM) module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from being generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

list scheduled-report

Displays all of the protocol security (PSM) scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics protocol-security report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics protocol-security report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics protocol-security report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics protocol-security scheduled-report(1)

analytics report

NAME

report - Displays an HTTP/L7-DoS analytics report.

MODULE
analytics

SYNTAX
Show, save or send an analytics report using the syntax shown in the following sections.

DISPLAY
show report view-by [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]
options:
drilldown {
 {
entity [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]
values
{
 [value ...]
}
 } ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
 [measure name ...]
}
order-by {
 {
 measure [measure name]
 sort-type [asc / desc]
 } ...
}
range [date range]

SAVE
save report view-by [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]
options:
drilldown {
 {
entity [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]
values
{
 [value ...]
}
 } ...
}
file [file name]
format [csv-aggregated | csv-time-series | pdf]
include-total
include-others
limit [number of rows]
measures {
 [measure name ...]
}
order-by {
 {
 measure [measure name]
 sort-type [asc / desc]
 } ...
}
range [date range]

SEND
send-mail report view-by [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]
options:
drilldown {
 {
entity [activity-type | application | attack-id | behavioral-signature | browser | client-ip | client-subnet |
country | country-code | device-id | dos-profile | http-method | http-transaction-outcome | mitigation |
os | pool-member | response-code | suspected-ip | trigger | url | user-agent | vector | virtual]
values
{
 [value ...]
}
 } ...
}
email-addresses {
 [email address ...]

```

}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate HTTP analytics reports. You can generate an HTTP analytics report for the following entities:

- Â· activity-type - Activity type.
- Â· application - Application services.
- Â· attack-id - Application/L7 DoS Attack ID.
- Â· behavioral-signature - Behavioral signature.
- Â· browser - Browser.
- Â· client-ip - A single client identified by an IP address.
- Â· client-subnet - Client subnet.
- Â· country - A country from which HTTP/HTTPS traffic was sent to each of the virtual servers.
- Â· country-code - Country code from which HTTP/HTTPS traffic was sent to each of the virtual servers.
- Â· device-id - Device ID.
- Â· dos-profile - DoS Profile.
- Â· http-method - Method.
- Â· http-transaction-outcome - HTTP Transaction outcomes (Blocked/Dropped/Passthrough/etc.)
- Â· mitigation - Mitigation.
- Â· os - OS name.
- Â· pool-member - Pool members.
- Â· response-code - An HTTP response code that was sent back to the client.
- Â· suspected-ip - Suspected address IP.
- Â· trigger - Trigger.
- Â· url - A URL accessed by HTTP or HTTPS.
- Â· user-agent - A browser identifier sent by the client's browser as part of the request for URL.
- Â· vector - Attack vector.
- Â· virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics report view-by virtual measures {average-tps} limit 20
```

Gets the average tps of 20 virtual servers (unordered).

```
show analytics report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } }
```

Gets the average tps of the top 20 virtual servers.

```
show analytics report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } range now-3d--now
```

Gets the average tps of the top 20 virtual servers from the last three days.

```
show analytics report view-by virtual drilldown { { entity application values { app } } { entity pool-member values { p1 p2 } } } range now-4d--now-2d measures {average-tps} limit 10 order-by { { measure average-tps sort-type DESC } }
```

Gets the average tps of the top 10 virtual servers (ordered by average tps) on app iApp (out of several monitored) on pool members p1 and p2 (out of five monitored p1-p5) in the interval ranging from two to four days ago.

```
show analytics report view-by response-code drilldown { { entity virtual values { v1 } } } measures { transactions }
```

Gets a distribution of requests per response code on virtual v1.

```
show analytics report view-by country drilldown { { entity application values { app } } } measures { average-concurrent-sessions average-sessions } order-by { { measure average-sessions sort-type DESC } } limit 5
```

Gets the new sessions and average concurrent sessions of the top five countries, ordered by the average concurrent sessions on the application app.

```
show analytics report view-by client-ip drilldown { { entity virtual values { v1 } } } measures { max-page-load-time } limit 1
```

Gets the client IP address with the worst page load time.

```
show analytics report view-by application drilldown { { entity pool-member values { p1 p2 } } } measures { transactions } order-by { { measure transactions } } range now-7d--now
```

Gets the distribution of requests per application on pool members p1 and p2 ordered by the number of requests during the last week.

```
save analytics report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format pdf file report.pdf
```

Gets the average tps of the top 20 virtual servers and exports to a PDF file on the BIG-IP system.

```
save analytics report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format csv-aggregated file report.csv
```

Gets the average tps of the top 20 virtual servers and exports to a CSV file on the BIG-IP system.

```
save analytics report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format csv-time-series file report.csv
```

Gets the average tps over time of the top 10 virtual servers and exports to a CSV file on the BIG-IP system.

```
send-mail analytics report view-by virtual measures {average-tps} limit 20 order-by { { measure average-tps sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

Gets the average tps over time of the top 10 virtual servers and sends out an email containing the report as a PDF.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The default value is transactions. The options are:

average-concurrent-sessions

The average number of concurrent sessions for each entity.

average-new-sessions
The average number of new sessions for each entity.

average-page-load-time
The average client page load time for each entity.

average-request-throughput
The average request throughput for each entity.

average-response-throughput
The average response throughput for each entity.

average-server-latency
The average server latency for each entity.

average-tps
The average number of transactions per second for each entity.

client-side-sampled-transactions
The number of transactions sampled for client side page load time.

max-page-load-time
The maximum client page load time for each entity.

max-request-throughput
The maximum request throughput for each entity.

max-response-throughput
The maximum response throughput for each entity.

max-server-latency
The maximum server latency for each entity.

max-tps
The maximum number of transactions per second for each entity.

transactions
The absolute number of transactions for each entity.

min-server-latency
The minimum server latency for each entity.

average-request-size
The average request size for each entity.

average-response-size
The average response size for each entity.

average-application-response-time
The average application response time for each entity.

min-application-response-time
The minimum application response time for each entity.

max-application-response-time
The maximum application response time for each entity.

average-client-ttfb
The average client TTFB for each entity.

min-client-ttfb
The minimum client TTFB for each entity.

max-client-ttfb
The maximum client TTFB for each entity.

average-clientside-network-latency
The average client-side network latency for each entity.

min-clientside-network-latency
The minimum client-side network latency for each entity.

max-clientside-network-latency
The maximum client-side network latency for each entity.

average-serverside-network-latency
The average server-side network latency for each entity.

min-serverside-network-latency
The minimum server-side network latency for each entity.

max-serverside-network-latency
The maximum server-side network latency for each entity.

average-request-duration
The average request duration for each entity.

`min-request-duration`
The minimum request duration for each entity.

`max-request-duration`
The maximum request duration for each entity.

`average-response-duration`
The average response duration for each entity.

`min-response-duration`
The minimum response duration for each entity.

`max-response-duration`
The maximum response duration for each entity.

`attacks-count`
The total number of attack for each entity.

`valid`
The total number of valid transactions for each entity.

`average-valid-tps`
The average number of valid transactions for each entity.

`mitigated`
The total number of mitigated transaction for each entity.

`average-mitigated-tps`
The average number of mitigated transaction for each entity.

`blocked`
The total number of blocked transactions for each entity.

`average-blocked-tps`
The average number of blocked transactions for each entity.

`incomplete`
The total number of incomplete transactions for each entity.

`average-incomplete-tps`
The average number of incomplete transactions for each entity.

`order-by`
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The default value for measures is previously chosen measures. The default value for sort type is desc (descending).

`range`
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

`smtp-config-override`
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

`show`, `save`, `send-mail`, `tmsh`, `lrm profile analytics`, `analytics report`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2018-07-05 analytics report(1)

analytics sip-dos report

NAME

`report` - Displays a SIP DoS analytics report.

MODULE

`analytics sip-dos`

SYNTAX

Show, save or send an analytics sip-dos report using the syntax shown in the following sections.

DISPLAY

```

show report view-by [ activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-me
options:
  drilldown {
  {
entity [ activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-method | sip-tran
values
{
  [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]

```

SAVE

```

save report view-by [ activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-met
options:
  drilldown {
  {
entity [ activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-method | sip-tran
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]

```

SEND

```

send-mail report view-by [ activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip
options:
  drilldown {
  {
entity [ activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-method | sip-tran
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate SIP DoS analytics reports. You can generate a SIP DoS prevention analytics report

for the following entities:

- Â· activity-type - Activity type.
- Â· application - Application services (iApp).
- Â· attack-id - DoS attack ID.
- Â· callee - Callee.
- Â· caller - Caller.
- Â· client-ip - Source IP Address.
- Â· country - Country.
- Â· country-code - Country code.
- Â· dos-profile - DoS profile.
- Â· mitigation - Mitigation.
- Â· sip-method - Method.
- Â· sip-transaction-outcome - Transaction outcome.
- Â· suspected-ip - Suspected IP Address.
- Â· trigger - Trigger.
- Â· vector - Attack vector.
- Â· virtual - Virtual server.
- Â· vlan - VLAN.

EXAMPLES

```
show analytics sip-dos report view-by attack-id
```

```
show analytics sip-dos report view-by vector drilldown { { entity method values { ACK } } }
```

```
send-mail analytics sip-dos report view-by callee limit 20 format pdf email-addresses {  
some.one@someaddress.com }
```

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. This option must be used with the drilldown option. You can also use it along with include-others.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

attacks-count

The total number of attacks for the selected view-by entity.

requests-count

The total number of requests that were received by the virtual server(/s)

requests-per-sec

The average number of requests that were received by the virtual server(/s)

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsb, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics sip-dos report(1)

analytics sip-dos scheduled-report

NAME

scheduled-report - Configure scheduled reports for SIP DoS.

MODULE

analytics sip-dos

SYNTAX

Configure the scheduled-report component within the analytics sip-dos module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |
replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE

delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the SIP DoS module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency  
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration

defined in `asm_smtp_conf`.

`modify scheduled-report myScheduledReport smtp-config none`

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from being generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to `person@domain.co` using the smtp configuration defined in `asm_smtp_conf`.

`list scheduled-report`

Displays all of the SIP DoS scheduled reports.

OPTIONS

`email-addresses`

A list of the email addresses of the recipients that receive the scheduled report.

`first-time`

First scheduled report time. Must be after current time and rounded up to the next round hour.

`frequency`

The scheduled report frequency. Example: `every-6-hours` means that the report will be generated and sent every 6 hours.

`include-total`

Enables or disables including a summary (Overall result) entity in results.

`multi-leveled-report`

Defines a custom multi-leveled report. Mutually exclusive with `predefined-report-name`. The multi-leveled-report definition contains the following parameters:

`chart-path`

A list of entities that define the scope in which the report will be displayed. For example: a chart path `{ violation url }` means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics sip-dos report.

`limit`

The number of view-by entities displayed in the scheduled report.

`time-diff`

The time range for the report.

`view-by`

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics sip-dos report.

`measures`

The measures which are available for the selected entities.

`predefined-report-name`

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with `multi-leveled-report`.

`smtp-config`

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

`device-group`

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

`list`, `modify`, `show`, `tmsh`, `analytics sip-dos report`, `sys smtp-server`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics sip-dos scheduled-report(1)

analytics sip report

NAME
report - Displays a SIP analytics report.

MODULE
analytics sip

SYNTAX
Show, save or send an analytics sip report using the syntax shown in the following sections.

DISPLAY
show report view-by [activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-me options:
drilldown {
 {
entity [activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-method | sip-tran values
{
 [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
 [measure name ...]
}
order-by {
 {
 measure [measure name]
 sort-type [asc | desc]
 } ...
}
range [date range]

SAVE
save report view-by [activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-met options:
drilldown {
 {
entity [activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-method | sip-tran values
{
 [value ...]
}
} ...
}
file [file name]
format [csv-aggregated | csv-time-series | pdf]
include-total
include-others
limit [number of rows]
measures {
 [measure name ...]
}
order-by {
 {
 measure [measure name]
 sort-type [asc | desc]
 } ...
}
range [date range]

SEND
send-mail report view-by [activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip options:
drilldown {
 {
entity [activity-type | application | attack-id | callee | caller | client-ip | country | country-code | dos-profile | mitigation | sip-method | sip-tran values
{
 [value ...]
}
} ...
}
email-addresses {
 [email address ...]
}
format [csv-aggregated | csv-time-series | pdf]
include-total
include-others
limit [number of rows]
measures {

```
[measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc | desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]
```

DESCRIPTION

Use this command to generate SIP DoS analytics reports. You can generate a SIP DoS prevention analytics report for the following entities:

- Â· activity-type - Activity type.
- Â· application - Application services (iApp).
- Â· attack-id - DoS attack ID.
- Â· callee - Callee.
- Â· caller - Caller.
- Â· client-ip - Source IP Address.
- Â· country - Country.
- Â· country-code - Country code.
- Â· dos-profile - DoS profile.
- Â· mitigation - Mitigation.
- Â· sip-method - Method.
- Â· sip-transaction-outcome - Transaction outcome.
- Â· suspected-ip - Suspected IP Address.
- Â· trigger - Trigger.
- Â· vector - Attack vector.
- Â· virtual - Virtual server.
- Â· vlan - VLAN.

EXAMPLES

```
show analytics sip report view-by attack-id
```

```
show analytics sip report view-by vector drilldown { { entity method values { ACK } } }
```

```
send-mail analytics sip report view-by callee limit 20 format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples see manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. This option must be used with the drilldown option. You can also use it along with include-others.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

attacks-count

The total number of attacks for the selected view-by entity.

requests-count

The total number of requests that were received by the virtual server(/s)

requests-per-sec

The average number of requests that were received by the virtual server(/s)

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2018-02-01 analytics sip report(1)

analytics sip scheduled-report

NAME

scheduled-report - Configure scheduled reports for SIP.

MODULE

analytics sip

SYNTAX

Configure the scheduled-report component within the analytics sip module using the syntax shown in the following sections.

CREATE/MODIFY

create scheduled-report [name]

modify scheduled-report [name]

options:

email-addresses [none | add | delete | modify |
replace-all-with] { email-address [string] }

first-time [date]

frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]

include-total [enabled | disabled]

multi-leveled-report {

chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }

limit [number of rows]

time-diff [last-hour | last-day | last-week | last-month | last-year]

view-by { entity name [string] }

measures [none | add | delete | modify | replace-all-with] { measure name [string] }

}

predefined-report-name [name]

smtp-config [name]

device-group [name]

DISPLAY

list scheduled-report

list scheduled-report [[[name] | [glob] | [regex]] ...]

show running-config scheduled-report

show running-config scheduled-report [[[name] | [glob] | [regex]] ...]

DELETE
delete scheduled-report [name]

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the SIP module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com } frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit 5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the SIP scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics sip report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics sip report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics sip report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

analytics ssl-orchestrator-service-virtual report

NAME

report - Displays a SSL Orchestrator Service Virtual Server Stats analytics report.

MODULE

analytics ssl-orchestrator-service-virtual

SYNTAX

Show, save or send a analytics ssl-orchestrator-service-virtual report using the syntax shown in the following sections.

DISPLAY

show report view-by [virtual]

options:

drilldown {

{

entity [virtual]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [virtual]

options:

drilldown {

{

entity [virtual]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [virtual]

options:

drilldown {

{

entity [virtual]

values

```

{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate SSL Orchestrator Service Virtual Server Stats analytics reports. You can generate a SSL Orchestrator Service Virtual Server Stats analytics report for the following entities:

- virtual - SSLO Service Virtual Servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics ssl-orchestrator-service-virtual report view-by virtual
```

```
show analytics ssl-orchestrator-service-virtual report view-by virtual drilldown { { entity virtual values {
virtual_1 virtual_2 } } }
```

```
send-mail analytics ssl-orchestrator-service-virtual report view-by virtual measures { server-side-bits-out }
limit 20 order-by { { measure server-side-bits-out sort-type desc } } format pdf email-addresses {
some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-server-side-concurrent-conns

Avg server side concurrent connections.

max-server-side-concurrent-conns

Max server side concurrent connections.

server-side-bits-in

Server side incoming bits.

server-side-bits-out
Server side outgoing bits.

server-side-conns
Server side connections.

order-by
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics ssl-orchestrator-service-virtual report(1)

analytics ssl-orchestrator-service-virtual scheduled-report

NAME

scheduled-report - Configure scheduled reports for SSL Orchestrator Service Virtual Server Stats.

MODULE

analytics ssl-orchestrator-service-virtual

SYNTAX

Configure the scheduled-report component within the analytics ssl-orchestrator-service-virtual module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the SSL Orchestrator Service Virtual Server Stats module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
```

```
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from being generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }  
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit  
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the SSL Orchestrator Service Virtual Server Stats scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics ssl-orchestrator-service-virtual report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics ssl-orchestrator-service-virtual report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics ssl-orchestrator-service-virtual report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2analytics ssl-orchestrator-service-virtual scheduled-report(1)

analytics ssl-orchestrator report

NAME

report - Displays a SSL Orchestrator analytics report.

MODULE

analytics ssl-orchestrator

SYNTAX

Show, save or send a analytics ssl-orchestrator report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ application | application-family | client-cipher-name | client-cipher-version | decryption-status |
  dest-country | ip-reputation | policy-action | server-cipher-name |
  server-cipher-version | service-path | site-ip | traffic-type | url-category |
  virtual ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ application | application-family | client-cipher-name | client-cipher-version | decryption-status |
  dest-country | ip-reputation | policy-action | server-cipher-name |
  server-cipher-version | service-path | site-ip | traffic-type | url-category |
  virtual ]
```

values

```
{
  [value ...]
```

```
}
```

```
} ...
```

```
}
```

```
field-fmt
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
  [measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
  measure [ measure name ]
```

```
  sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SAVE

```
save report view-by [ application | application-family | client-cipher-name | client-cipher-version | decryption-status |
  dest-country | ip-reputation | policy-action | server-cipher-name |
  server-cipher-version | service-path | site-ip | traffic-type | url-category |
  virtual ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ application | application-family | client-cipher-name | client-cipher-version | decryption-status |
  dest-country | ip-reputation | policy-action | server-cipher-name |
  server-cipher-version | service-path | site-ip | traffic-type | url-category |
  virtual ]
```

values

```
{
  [value ...]
```

```
}
```

```
} ...
```

```
}
```

```
file [ file name ]
```

```
format [ csv-aggregated | csv-time-series | pdf ]
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
  [measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
  measure [ measure name ]
```

```
  sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SEND

```
send-mail report view-by [ application | application-family | client-cipher-name | client-cipher-version | decryption-status |
  dest-country | ip-reputation | policy-action | server-cipher-name |
  server-cipher-version | service-path | site-ip | traffic-type | url-category |
  virtual ]
```

```

options:
drilldown {
{
entity [ application | application-family | client-cipher-name | client-cipher-version | decryption-status |
dest-country | ip-reputation | policy-action | server-cipher-name |
server-cipher-version | service-path | site-ip | traffic-type | url-category |
virtual ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate SSL Orchestrator analytics reports. You can generate a SSL Orchestrator analytics report for the following entities:

- application - Application Name.
- application-family - Application Family.
- client-cipher-name - Client Cipher Name.
- client-cipher-version - Client Cipher Version.
- decryption-status - Decryption status.
- dest-country - Destination Country.
- ip-reputation - IP reputation.
- policy-action - Policy action.
- server-cipher-name - Server Cipher Name.
- server-cipher-version - Server Cipher Version.
- service-path - Service Path.
- site-ip - Site IP.
- traffic-type - Traffic type.
- url-category - URL category.
- virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics ssl-orchestrator report view-by application
```

```
show analytics ssl-orchestrator report view-by application drilldown { { entity virtual values { virtual_1
virtual_2 } } }
```

```
send-mail analytics ssl-orchestrator report view-by application measures { hits-count-per-sec } limit 20
order-by { { measure hits-count-per-sec sort-type desc } } format pdf email-addresses {
some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device
Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list
Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

client-bytes-in

Client Bytes In.

client-bytes-in-per-sec

Client Bytes In Per Sec.

client-bytes-out

Client Bytes Out.

client-bytes-per-sec

Client Bytes Out Per Second.

duration

Duration.

hits-count

Hits Count.

hits-count-per-sec

Hits Count Per Second.

server-bytes-in

Server Bytes In.

server-bytes-in-per-sec

Server Bytes In Per Sec.

server-bytes-out

Server Bytes Out.

server-bytes-out-per-sec

Server Bytes Out Per Second.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

analytics ssl-orchestrator scheduled-report

NAME

scheduled-report - Configure scheduled reports for SSL orchestrator.

MODULE

analytics ssl-orchestrator

SYNTAX

Configure the scheduled-report component within the analytics ssl-orchestrator module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the SSL orchestrator module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the SSL orchestrator scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent

every 6 hours.

`include-total`

Enables or disables including a summary (Overall result) entity in results.

`multi-leveled-report`

Defines a custom multi-leveled report. Mutually exclusive with `predefined-report-name`. The multi-leveled-report definition contains the following parameters:

`chart-path`

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics ssl-orchestrator report.

`limit`

The number of view-by entities displayed in the scheduled report.

`time-diff`

The time range for the report.

`view-by`

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics ssl-orchestrator report.

`measures`

The measures which are available for the selected entities.

`predefined-report-name`

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with `multi-leveled-report`.

`smtp-config`

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

`device-group`

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

`list`, `modify`, `show`, `tmsh`, `analytics ssl-orchestrator report`, `sys smtp-server`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics ssl-orchestrator scheduled-report(1)

analytics swg-blocked report

NAME

report - Displays an swg-blocked analytics report.

MODULE

analytics swg-blocked

SYNTAX

Show, save or send an analytics swg-blocked report using the syntax shown in the following sections.

DISPLAY

show report view-by [ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | security-category | host-name] options:

drilldown {

{ entity [ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | security-category | host-name] values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

```

measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SAVE

```

save report view-by [ ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | security-category | host-name ]
options:
  drilldown {
    {
entity [ ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | security-category | host-name ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

```

SEND

```

send-mail report view-by [ ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | security-category | host-name ]
options:
  drilldown {
    {
entity [ ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | security-category | host-name ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate swg-blocked analytics reports. You can generate an application-security-network analytics report for the following entities:

- Â· ssl_bypass - Is HTTP/HTTPS inspection bypassed
- Â· username - User name
- Â· client-ip - Client IP
- Â· host-name - Host name
- Â· url - URL
- Â· category - URL category
- Â· security-category - Categories which their parent is 'Security' category

• url-filter - URL filter

• scheme - Scheme

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics swg-blocked report view-by client-ip
```

```
show analytics swg-blocked report view-by client-ip drilldown { { entity scheme values { my_scheme } } }
```

```
send-mail analytics swg-blocked report view-by category measures { blocked-count } limit 20 order-by { { measure blocked-count sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

blocked-count

The total number of blocked requests for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2016-01-07 analytics swg-blocked report(1)

analytics swg-blocked scheduled-report

NAME

scheduled-report - Configure scheduled reports for SWG (blocked).

MODULE

analytics swg-blocked

SYNTAX

Configure the scheduled-report component within the analytics swg-blocked module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the SWG (blocked) module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the SWG (blocked) scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics swg-blocked report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics swg-blocked report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics swg-blocked report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics swg-blocked scheduled-report(1)

analytics swg report

NAME

report - Displays an swg analytics report.

MODULE

analytics swg

SYNTAX

Show, save or send an analytics swg report using the syntax shown in the following sections.

DISPLAY

show report view-by [action | ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | host-name]

options:

drilldown {

{

entity [action | ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | host-name]

values

{

[value ...]

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

```

}
range [date range]

SAVE
save report view-by [ action | ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | host-name ]
options:
  drilldown {
  {
entity [ action | ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | host-name ]
values
{
[value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]

SEND
send-mail report view-by [ action | ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | host-name ]
options:
  drilldown {
  {
entity [ action | ssl_bypass | username | client_ip | url | category | url_filter | filter_policy | host-name ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate swg analytics reports. You can generate an application-security-network analytics report for the following entities:

- Â· action - Transaction's action (Blocked / Allowed / Allowed with log)
- Â· ssl_bypass - Is HTTP/HTTPS inspection bypassed
- Â· username - User name
- Â· client-ip - Client IP
- Â· host-name - Host name
- Â· url - URL
- Â· category - URL category
- Â· url-filter - URL filter
- Â· scheme - Scheme

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

show analytics swg report view-by client-ip

show analytics swg report view-by client-ip drilldown { { entity scheme values { my_scheme } } }

send-mail analytics swg report view-by category measures { blocked-count } limit 20 order-by { { measure blocked-count sort-type desc } } format pdf email-addresses { some.one@someaddress.com }

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

request-count

The total number of requests for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2016-01-07 analytics swg report(1)

analytics swg scheduled-report

NAME

scheduled-report - Configure scheduled reports for SWG.

MODULE

analytics swg

SYNTAX

Configure the scheduled-report component within the analytics swg module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the SWG module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the SWG scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics swg report.

limit

The number of view-by entities displayed in the scheduled report.

`time-diff`
The time range for the report.

`view-by`
The main entity that the report is viewed by. For a list of valid entities see the help manual for `analytics swg report`.

`measures`
The measures which are available for the selected entities.

`predefined-report-name`
Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with `multi-leveled-report`.

`smtp-config`
Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

`device-group`
Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

`list`, `modify`, `show`, `tmsh`, `analytics swg report`, `sys smtp-server`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 `analytics swg scheduled-report(1)`

analytics system-monitor report

NAME
`report` - Displays a System Monitor analytics report.

MODULE
`analytics system-monitor`

SYNTAX
Show, save or send a analytics system-monitor report using the syntax shown in the following sections.

DISPLAY
`show report view-by [slot]`
options:
`drilldown {`
 `{`
 entity [slot]
 values
 `{`
 [value ...]
 `}`
 } ...
 `}`
`field-fmt`
`include-total`
`include-others`
`limit [number of rows]`
`measures {`
 [measure name ...]
`}`
`order-by {`
 `{`
 measure [measure name]
 sort-type [asc / desc]
 } ...
`}`
`range [date range]`

SAVE
`save report view-by [slot]`
options:
`drilldown {`
 `{`

```

entity [ slot ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ slot ]
options:
drilldown {
  {
entity [ slot ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate System Monitor analytics reports. You can generate a System Monitor analytics report for the following entities:

• slot - Slot ID.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics system-monitor report view-by slot
```

```
show analytics system-monitor report view-by slot drilldown { { entity slot values { 0 1 } } }
```

```
send-mail analytics system-monitor report view-by slot measures { avg-cpu } limit 20 order-by { { measure avg-cpu sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

`format`

Specifies the exported file format to be saved or sent. This option must be specified when using the `save` or `send-mail` commands.

`include-others`

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with `include-total`.

`include-total`

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

`limit`

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

`measures`

Specifies a list of measures that can be used with the chosen entity type. The options are:

`avg-concurrent-connections`

An average number of concurrent connections at any given time.

`avg-cpu`

Average TMM CPU usage over the selected time period.

`avg-memory`

Average RAM usage over the selected time period.

`avg-throughput`

Average bidirectional client throughput.

`concurrent-connections-health`

Concurrent connections health percent.

`cpu-health`

CPU health percent.

`max-concurrent-connections`

Concurrent Connections Max value.

`max-cpu`

TMM CPU Max value.

`max-memory`

Memory Max value.

`max-throughput`

Throughput Max value.

`memory-health`

Memory health percent.

`throughput-health`

Throughput health percent.

`order-by`

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is `desc` (descending).

`range`

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (`now--now-1h`).

`smtp-config-override`

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

`show`, `save`, `send-mail`, `tms`, `itm profile analytics`, `analytics`, `analytics report`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

analytics tcp-analytics report

NAME

report - Displays the TCP analytics report.

MODULE

analytics tcp analytics

SYNTAX

Show, save or send an analytics tcp-analytics report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ application | city-postcode | continent-code |
country-region | nexthop | remote-address | request-side | subnet | user-key | virtual ]
options:
  drilldown {
  {
  entity [ application | city-postcode | continent-code | country-region | nexthop | remote-address | request-side | subnet | user-key | virtual ]
  values
  {
  [value ...]
  }
  } ...
  }
  field-fmt
  include-total
  include-others
  limit [number of rows]
  measures {
  [measure name ...]
  }
  order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
  }
  range [date range]
```

SAVE

```
save report view-by [ application | city-postcode | continent-code |
country-region | nexthop | remote-address | request-side | subnet | user-key | virtual ]
options:
  drilldown {
  {
  entity [ application | city-postcode | continent-code | country-region | nexthop | remote-address | request-side | subnet | user-key | virtual ]
  values
  {
  [value ...]
  }
  } ...
  }
  file [ file name ]
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
  [measure name ...]
  }
  order-by {
  {
  measure [ measure name ]
  sort-type [ asc / desc ]
  } ...
  }
  range [date range]
```

SEND

```
send-mail report view-by [ application | city-postcode | continent-code |
country-region | nexthop | remote-address | request-side | subnet | user-key | virtual ]
options:
  drilldown {
  {
  entity [ application | city-postcode | continent-code | country-region | nexthop | remote-address | request-side | subnet | user-key | virtual ]
  values
  {
  [value ...]
  }
  } ...
  }
  email-addresses {
  [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
```

```

include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate tcp analytics reports. You can generate a TCP analytics report for the following entities:

- Â· application - Application services
- Â· city-postcode - City and postcode of remote host
- Â· continent-code - Continent of remote host
- Â· country-region - Country and region of remote host
- Â· nexthop - MAC address of next routing hop
- Â· remote-address - IP address of remote host
- Â· request-side - Client-side or server-side connection
- Â· subnet - /24 subnet of remote host
- Â· user-key - Key user-defined by TCP::analytics iRule.
- Â· virtual - Virtual Server

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics tcp-analytics report view-by virtual
```

```
show analytics tcp-analytics report view-by virtual drilldown { { entity remote-address values { 10.10.2.2 } } }
```

```
send-mail analytics tcp-analytics report view-by remote-address measures { goodput-rcv } limit 20 order-by { { measure goodput-rcv sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-connection-goodput-received

The mean incoming goodput (in bits per second) of all TCP connections for the selected filter (entity).

avg-connection-goodput-sent
The mean outgoing goodput (in bits per second) of all TCP connections for the selected filter (entity).

connections-closed
The total number of connections ended for the selected filter (entity).

connections-length-mean
The mean connection length (in ms) for the selected filter (entity).

connections-opened
The total number connections started for the selected filter (entity).

delaystate-3whs
Time (in ms) spent by the connection in delay state three-way handshake (3WHS) for the selected filter (entity).

delaystate-app
Time (in ms) spent by the connection in delay state limited by lack of application data for the selected filter (entity).

delaystate-closing
Time (in ms) spent by the connection in delay state waiting for ack of FIN for the selected filter (entity).

delaystate-cwnd
Time (in ms) spent by the connection in delay state limited by the congestion window for the selected filter (entity).

delaystate-nagle
Time (in ms) spent by the connection in delay state where the app is limited, last packet held due to Nagle algorithm for the selected filter (entity).

delaystate-ratepace
Time (in ms) spent by the connection in delay state where the app is limited, transmission delayed by rate pacing for the selected filter (entity).

delaystate-retx
Time (in ms) spent by the connection in delay state retransmitting lost packets for the selected filter (entity).

delaystate-rwnd
Time (in ms) spent by the connection in delay state where the client receive window is limited for the selected filter (entity).

delaystate-sndbuf
Time (in ms) spent by the connection in delay state limited by send buffer configuration for the selected filter (entity).

delaystate-waitforack
Time (in ms) spent by the connection in delay state waiting for acknowledgement for the selected filter (entity).

goodput-received
The aggregate data rate (in bits per second) received over all connections for the selected filter (entity).

goodput-sent
The aggregate data rate (in bits per second) sent over all connections for the selected filter (entity).

max-connection-goodput-received
The maximum incoming goodput (in bits per second) on any TCP connection for the selected filter (entity).

max-connection-goodput-sent
The maximum outgoing goodput (in bits per second) on any TCP connection for the selected filter (entity).

min-connection-goodput-received
The minimum incoming goodput (in bits per second) on any TCP connection for the selected filter (entity).

min-connection-goodput-sent
The minimum outgoing goodput (in bits per second) on any TCP connection for the selected filter (entity).

packet-loss-rate
The fraction of sent packets that are lost for the selected filter (entity).

packets-lost
The total number of sent packets lost for the selected filter (entity).

packets-received

The total number of packets received for the selected filter (entity).

packets-sent
The total number of packets sent for the selected filter (entity).

rtt-avg
The mean round trip time (in ms) for the selected filter (entity).

rtt-max
The maximum round trip time in ms for the selected filter (entity).

rtt-min
The minimum round trip time (in ms) for the selected filter (entity).

rttvar-mean
The mean variance of the round trip time (in ms) for the selected filter (entity).

SEE ALSO

show, save, send-mail, tmsb, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2017-06-14 analytics tcp-analytics report(1)

analytics tcp-analytics scheduled-report

NAME

scheduled-report - Configure scheduled reports for TCP analytics.

MODULE

analytics tcp-analytics

SYNTAX

Configure the scheduled-report component within the analytics tcp-analytics module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the TCP analytics module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration

defined in `asm_smtp_conf`.

`modify scheduled-report myScheduledReport smtp-config none`

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to `person@domain.co` using the smtp configuration defined in `asm_smtp_conf`.

`list scheduled-report`

Displays all of the TCP analytics scheduled reports.

OPTIONS

`email-addresses`

A list of the email addresses of the recipients that receive the scheduled report.

`first-time`

First scheduled report time. Must be after current time and rounded up to the next round hour.

`frequency`

The scheduled report frequency. Example: `every-6-hours` means that the report will be generated and sent every 6 hours.

`include-total`

Enables or disables including a summary (Overall result) entity in results.

`multi-leveled-report`

Defines a custom multi-leveled report. Mutually exclusive with `predefined-report-name`. The multi-leveled-report definition contains the following parameters:

`chart-path`

A list of entities that define the scope in which the report will be displayed. For example: a chart path `{ violation url }` means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics `tcp-analytics report`.

`limit`

The number of view-by entities displayed in the scheduled report.

`time-diff`

The time range for the report.

`view-by`

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics `tcp-analytics report`.

`measures`

The measures which are available for the selected entities.

`predefined-report-name`

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with `multi-leveled-report`.

`smtp-config`

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

`device-group`

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

`list`, `modify`, `show`, `tmsh`, `analytics tcp-analytics report`, `sys smtp-server`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics tcp-analytics scheduled-report(1)

analytics tcp report

NAME

report - Displays an tcp analytics report.

MODULE

analytics tcp

SYNTAX

Show, save or send an analytics tcp report using the syntax shown in the following sections.

DISPLAY

show report view-by [virtual | tcp]

options:

drilldown {

{

entity [virtual | tcp]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [virtual | tcp]

options:

drilldown {

{

entity [virtual | tcp]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SEND

send-mail report view-by [virtual | tcp]

options:

drilldown {

{

entity [virtual | tcp]

values

{

[value ...]

}

} ...

}

email-addresses {

[email address ...]

}

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

```

[measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate tcp analytics reports. You can generate a TCP analytics report for the following entities:

• virtual - Virtual Server

• tcp - TCP Profile

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics tcp report view-by virtual
```

```
show analytics tcp report view-by virtual drilldown { { entity virtual values { 172.12.34.56 } } }
```

```
send-mail analytics tcp report view-by tcp measures { max-active-conns } limit 20 order-by { { measure max-active-conns sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

avg-active-conns

The average number of active connections for the selected filter (entity).

max-active-conns

The max number of active connections for the selected filter (entity).

total-accepts

The total number of accepted connections for the selected filter (entity).

total-accept_fails

The total number of denied accept connections for the selected filter (entity).

total-new-conns

The total number of new connections for the selected filter (entity).

total-failed-conns

The total number of failed connections for the selected filter (entity).

total-expired-conns

The total number of expired connections for the selected filter (entity).

total-abandoned-connections
The total number of abandoned connections for the selected filter (entity).

total-rst-packets
The total number of RST connections for the selected filter (entity).

total-malformed-segments
The total number of malformed connections for the selected filter (entity).

total-oo-segs
The total number of out of ordered segments for the selected filter (entity).

total-rx-cookie
The total number of received SYN cookies for the selected filter (entity).

total-rxbadcookies
The total number of received bad SYN cookies for the selected filter (entity).

total-hw-cookies
The total number of received HW SYN cookies for the selected filter (entity).

total-syncacheover
The total number of SYN cache overflow for the selected filter (entity).

total-txrextmits
The total number of retransmitted segments for the selected filter (entity).

total-sndpack
The total number of sent packets for the selected filter (entity).

order-by
Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range
Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO
show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-07-10 analytics tcp report(1)

analytics tcp scheduled-report

NAME
scheduled-report - Configure scheduled reports for TCP.

MODULE
analytics tcp

SYNTAX
Configure the scheduled-report component within the analytics tcp module using the syntax shown in the following sections.

CREATE/MODIFY
create scheduled-report [name]
modify scheduled-report [name]
options:
email-addresses [none | add | delete | modify |
replace-all-with] { email-address [string] }
first-time [date]
frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
include-total [enabled | disabled]
multi-leveled-report {
chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
limit [number of rows]

```
time-diff [last-hour | last-day | last-week | last-month | last-year]
view-by { entity name [string] }
measures [none | add | delete | modify | replace-all-with] { measure name [string] }
}
predefined-report-name [name]
smtp-config [name]
device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the TCP module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the TCP scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics tcp report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics tcp report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics tcp report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics tcp scheduled-report(1)

analytics tmm-dns-zone report

NAME

report - Displays a TMM DNS Zone analytics report.

MODULE

analytics tmm-dns-zone

SYNTAX

Show, save or send a analytics tmm-dns-zone report using the syntax shown in the following sections.

DISPLAY

show report view-by [zone]

options:

drilldown {

{

entity [zone]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [zone]

options:

drilldown {

{

entity [zone]

values

{

[value ...]

}

} ...

}

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

```

    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ zone ]
options:
drilldown {
  {
entity [ zone ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
  {
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate TMM DNS Zone analytics reports. You can generate a TMM DNS Zone analytics report for the following entities:

- zone - TMM DNS zone.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics tmm-dns-zone report view-by zone
```

```
show analytics tmm-dns-zone report view-by zone drilldown { { entity zone values { value } } }
```

```
send-mail analytics tmm-dns-zone report view-by zone measures { responses } limit 20 order-by { { measure responses sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

a-queries

Dnsx qtype IPV4 address.

aaaa-queries

Dnsx qtype IPV6 address.

any-queries

Dnsx qType any.

axfr-queries

Dnsx qtype axfr.

cname-queries

Dnsx qtype cname.

ixfr-queries

Dnsx qtype ixfr.

mx-queries

Dnsx qtype mx.

ns-queries

Dnsx qtype ns.

other-queries

Dnsx qtype other.

queries

Dnsx queries.

responses

Dnsx response.

soa-queries

Dnsx qtype soa.

srv-queries

Dnsx qtype srv.

tsig-bad-key

Transaction signature bad key.

tsig-bad-sig

Transaction signature bad sig.

tsig-bad-time

Transaction signature bad time.

tsig-missing

Transaction signature missing.

tsig-not-required

Transaction signature not required.

tsig-verified

Transaction signature verified.

txt-queries

Dnsx qtype txt.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

analytics traffic-classification report

NAME

report - Displays a Traffic Classification analytics report.

MODULE

analytics traffic-classification

SYNTAX

Show, save or send a analytics traffic-classification report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ action-name | application | application-category | destination-country | destination-ip | destination-port |
  profiles | risk | source-country | source-ip | url-category | user-name |
  virtual ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ action-name | application | application-category | destination-country | destination-ip | destination-port |
  profiles | risk | source-country | source-ip | url-category | user-name |
  virtual ]
```

values

```
{
  [value ...]
}
```

```
} ...
}
```

field-fmt

include-total

include-others

limit [number of rows]

measures {

```
[measure name ...]
```

```
}
```

order-by {

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

range [date range]

SAVE

```
save report view-by [ action-name | application | application-category | destination-country | destination-ip | destination-port |
  profiles | risk | source-country | source-ip | url-category | user-name |
  virtual ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ action-name | application | application-category | destination-country | destination-ip | destination-port |
  profiles | risk | source-country | source-ip | url-category | user-name |
  virtual ]
```

values

```
{
  [value ...]
}
```

```
} ...
}
```

file [file name]

format [csv-aggregated | csv-time-series | pdf]

include-total

include-others

limit [number of rows]

measures {

```
[measure name ...]
```

```
}
```

order-by {

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

range [date range]

SEND

```
send-mail report view-by [ action-name | application | application-category | destination-country | destination-ip | destination-port |
  profiles | risk | source-country | source-ip | url-category | user-name |
  virtual ]
```

options:

```

drilldown {
  {
entity [ action-name | application | application-category | destination-country | destination-ip | destination-port |
        profiles | risk | source-country | source-ip | url-category | user-name |
        virtual ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate Traffic Classification analytics reports. You can generate a Traffic Classification analytics report for the following entities:

- Â· action-name - Actions Name.
- Â· application - Classification application Name.
- Â· application-category - Application classification category.
- Â· destination-country - Destination countries.
- Â· destination-ip - A single destination identified by an IP address.
- Â· destination-port - Destination ports.
- Â· profiles - Classification profile.
- Â· risk - Risk.
- Â· source-country - Country from which the traffic originated.
- Â· source-ip - A single source identified by an IP address.
- Â· url-category - URL category.
- Â· user-name - Users name.
- Â· virtual - Virtual servers.

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics traffic-classification report view-by action-name
```

```
show analytics traffic-classification report view-by action-name drilldown { { entity application values {
application_name } } }
```

```
send-mail analytics traffic-classification report view-by action-name measures { total-bytes-in } limit 20
order-by { { measure total-bytes-in sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple

(not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

average-throughput

Total average throughput.

average-throughput-in

Average incoming throughput.

average-throughput-out

Average outgoing throughput.

flows

Total number of classified flows.

total-bytes

Total bytes in both incoming and outgoing traffic.

total-bytes-in

Byte counter of incoming (uplink) traffic.

total-bytes-out

Byte counter of outgoing (downlink) traffic.

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-07-04 analytics traffic-classification report(1)

analytics traffic-classification scheduled-report

NAME

scheduled-report - Configure scheduled reports for traffic classification.

MODULE

analytics traffic-classification

SYNTAX

Configure the scheduled-report component within the analytics traffic-classification module using the syntax shown in the following sections.

```
CREATE/MODIFY
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
limit [number of rows]
time-diff [last-hour | last-day | last-week | last-month | last-year]
view-by { entity name [string] }
measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

```
DISPLAY
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

```
DELETE
delete scheduled-report [name]
```

DESCRIPTION
Use the scheduled-report component to create, modify or delete scheduled reports for the traffic classification module.

EXAMPLES
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the traffic classification scheduled reports.

OPTIONS
email-addresses
A list of the email addresses of the recipients that receive the scheduled report.

first-time
First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency
The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total
Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report
Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path
A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics traffic-classification report.

limit
The number of view-by entities displayed in the scheduled report.

time-diff
The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics traffic-classification report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics traffic-classification report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics traffic-classification scheduled-report(1)

analytics udp report

NAME

report - Displays an udp analytics report.

MODULE

analytics udp

SYNTAX

Show, save or send an analytics udp report using the syntax shown in the following sections.

DISPLAY

```
show report view-by [ virtual | udp ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ virtual | udp ]
```

```
values
```

```
{
```

```
[value ...]
```

```
}
```

```
} ...
```

```
}
```

```
field-fmt
```

```
include-total
```

```
include-others
```

```
limit [number of rows]
```

```
measures {
```

```
[measure name ...]
```

```
}
```

```
order-by {
```

```
{
```

```
measure [ measure name ]
```

```
sort-type [ asc / desc ]
```

```
} ...
```

```
}
```

```
range [date range]
```

SAVE

```
save report view-by [ virtual | udp ]
```

options:

```
drilldown {
```

```
{
```

```
entity [ virtual | udp ]
```

```
values
```

```
{
```

```
[value ...]
```

```

}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]

SEND
send-mail report view-by [ virtual | udp ]
options:
  drilldown {
    {
  entity [ virtual | udp ]
  values
  {
    [value ...]
  } ...
  }
  email-addresses {
    [email address ...]
  }
  format [ csv-aggregated | csv-time-series | pdf ]
  include-total
  include-others
  limit [number of rows]
  measures {
    [measure name ...]
  }
  order-by {
    {
      measure [ measure name ]
      sort-type [ asc / desc ]
    } ...
  }
  range [date range]
  smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate udp analytics reports. You can generate a UDP analytics report for the following entities:

• virtual - Virtual Server

• udp - UDP Profile

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics udp report view-by virtual
```

```
show analytics udp report view-by virtual drilldown { { entity virtual values { 172.12.34.56 } } }
```

```
send-mail analytics udp report view-by udp measures { max-active-conns } limit 20 order-by { { measure max-active-conns sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

`include-others`

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

`include-total`

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

`limit`

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

`measures`

Specifies a list of measures that can be used with the chosen entity type. The options are:

`avg-active-conns`

The average number of active connections for the selected filter (entity).

`max-active-conns`

The max number of active connections for the selected filter (entity).

`total-accepts`

The total number of accepted connections for the selected filter (entity).

`total-accept_fails`

The total number of denied accept connections for the selected filter (entity).

`total-new-conns`

The total number of new connections for the selected filter (entity).

`total-failed-conns`

The total number of failed connections for the selected filter (entity).

`total-expired-conns`

The total number of expired connections for the selected filter (entity).

`total-received-datagrams`

The total number of received datagrams for the selected filter (entity).

`total-malformed-datagrams`

The total number of malformed datagrams for the selected filter (entity).

`total-icmp-unreachable`

The total number of ICMP unreachable for the selected filter (entity).

`total-bad-sum-datagrams`

The total number of bad checksum datagrams for the selected filter (entity).

`total-no-sum-datagrams`

The total number of 'no checksum' datagrams for the selected filter (entity).

`total-transmitted-datagrams`

The total number of transmitted datagrams for the selected filter (entity).

`order-by`

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

`range`

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

`smtp-config-override`

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tms, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

analytics udp scheduled-report

NAME

scheduled-report - Configure scheduled reports for UDP.

MODULE

analytics udp

SYNTAX

Configure the scheduled-report component within the analytics udp module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the UDP module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the UDP scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics udp report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics udp report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tms, analytics udp report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics udp scheduled-report(1)

analytics uri-type

NAME

uri-type - Configure uri-type

MODULE

analytics uri-type

SYNTAX

Configure the uri-types components within the analytics module using the syntax shown in the following sections.

CREATE/MODIFY

```
create uri-type [name]
modify uri-type [name]
properties:
  file-extensions
  [add | delete | replace-all-with] {
name [string]
options:
  name - file extension name
}
```

DISPLAY

```
list uri-type
list uri-type [ [ [name] | [glob] | [regex] ] ...]
options:
  all
  all-properties
  one-line
```

properties:

file-extensions

DELETE

delete uri-type [[name] | all]

DESCRIPTION

Use the analytics uri-type command to create / list / modify uri-types.

uri-type is a definitions for URIs that should be collected and reported under predefined name value.

Each uri-type has a name value and a file-extensions list associated with it.

Once uri-type(s) created it will be associated with all AVR profiles (see below), any URI that has a match, will be reported as uri-type name value under view-by url analytics report. For example: create /analytics uri-type Images file-extensions add { png jpg gif ico } /images/logo.png is a URI match png and will be reported as "Image" file extension. /favicon.ico is a URI match ico and will be reported as "Image" file extension. /images/pic.jpg is a URI match jpg and will be reported as "Image" file extension.

Limitations may apply on total number of uri-types object and total number of file extensions associated with them, cross the whole system.

uri-type can only be defined under partition /Common.

Notes:

uri-type name is unique - there can't be two uri-types with the same name value.

uri-type file-extensions list is unique - any file extension can be defined only once. Means, it is not possible to share same file extensions across multiple uri-type configuration objects.

uri-type is a global analytics configuration. All AVR profiles will be affected by every configured uri-type object.

name The identifier of the uri-type and the actual URL replace value. This value will be seen in analytics reports instead of the original URI. Eg: In create uri-type Image file-extensions add { jpg png gif } example, the name is Image and file-extensions are: jpg, png, gif.

file-extensions

List of at least one or more file extensions (aka file types / file suffixes) that belongs to this uri-type. The value is provided in lower case characters but it represents case insensitive match. Eg: 'jpg' will match all of the following file extensions: JPG, jPg, jpG, JPg etc...

EXAMPLES

Display all current uri-types:

list uri-types

Create new uri-type to associate and report certain images files as "Image" URL:

create uri-type Images file-extensions add { jpg png gif ico icon jjpg }

Modify uri-type, associate more file extensions with it:

modify uri-type Images file-extensions add { jpeg }

Modify uri-type, remove and disassociate file extension:

modify uri-type Images file-extensions delete { jjpg }

Delete uri-type:

delete uri-type Images

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-06-27 analytics uri-type(1)

analytics vcmp report

NAME

report - Displays an vcmp analytics report.

MODULE

analytics vcmp

SYNTAX

Show, save or send an analytics vcmp report using the syntax shown in the following sections.

DISPLAY

show report view-by [guest | slot | interface | process-name]

options:

```
drilldown {
  {
entity [ guest | slot | interface | process-name ]
values
{
  [value ...]
}
} ...
}
field-fmt
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SAVE

save report view-by [guest | slot | interface | process-name]

options:

```
drilldown {
  {
entity [ guest | slot | interface | process-name ]
values
{
  [value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
range [date range]
```

SEND

send-mail report view-by [guest | slot | interface | process-name]

options:

```
drilldown {
  {
entity [ guest | slot | interface | process-name ]
values
{
  [value ...]
}
} ...
}
email-addresses {
  [email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
  [measure name ...]
}
order-by {
  {
    measure [ measure name ]
    sort-type [ asc / desc ]
  } ...
}
}
```

range [date range]
smtp-config-override [smtp configuration object name]

DESCRIPTION

Use this command to generate vcmp analytics reports. You can generate a vCMP analytics report for the following entities:

Â· slot - Slot ID

Â· guest - vCMP Guest name

Â· interface - vCMP Interface

Â· process-name - Process name

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

show analytics vcmp report view-by slot

show analytics vcmp report view-by slot drilldown { { entity slot values { 5 } } }

send-mail analytics vcmp report view-by guest measures { network-bytes-in } limit 20 order-by { { measure network-bytes-in sort-type desc } } format pdf email-addresses { some.one@someaddress.com }

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple (not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

network-bytes-in

The total number of external and internal bytes received for the selected filter (entity).

network-bytes-out

The total number of external and internal bytes sent for the selected filter (entity).

network-average-in-throughput

The average internal and external received throughput for the selected filter (entity).

network-average-out-throughput

The average internal and external throughput sent for the selected filter (entity).

average-guest-cpu-usage

The average number of CPU usage for the selected filter (entity) relative to host CPU.

diskio-bytes-read

The average number of bytes read from disk for the selected filter (entity).

diskio-bytes-written

The average number of bytes written to disk for the selected filter (entity).

diskio-requests-read

The average number of read requests from disk for the selected filter (entity).

diskio-requests-written

The average number of write requests from disk for the selected filter (entity).

average-process-cpu-usage

The average number of CPU usage for the selected filter (entity) normalized relatively to guest CPU.

traffic-client-new-connections

The total number of client-side new connections for the selected filter (entity).

traffic-client-avg-connections

The average number of client-side connections opened for a selected filter (entity).

traffic-client-packets-in

The total number of client-side received packets for a selected filter (entity).

traffic-client-packets-out

The total number of client-side sent packets for a selected filter (entity).

traffic-client-bytes-in

The total number of client-side received bytes for a selected filter (entity).

traffic-client-bytes-out

The total number of client-side sent bytes for a selected filter (entity).

traffic-server-new-connections

The total number of server-side new connections for the selected filter (entity).

traffic-server-avg-connections

The average number of server-side connections opened for a selected filter (entity).

traffic-server-packets-in

The total number of server-side received packets for a selected filter (entity).

traffic-server-packets-out

The total number of server-side sent packets for a selected filter (entity).

traffic-server-bytes-in

The total number of server-side received bytes for a selected filter (entity).

traffic-server-bytes-out

The total number of server-side sent bytes for a selected filter (entity).

total-assisted-connections

The total number of all hardware accelerated assisted connections for a selected filter (entity).

current-assisted-connections

The average number of all hardware accelerated assisted connections for a selected filter (entity).

hardware-syncookies-generated

The total number of SYN cookies generated for a selected filter (entity).

hardware-syncookies-detected

The total number of SYN cookies detected for a selected filter (entity).

hw-accel-client-packets-in

The total number of hardware accelerated client-side received packets for a selected filter (entity).

hw-accel-client-packets-out

The total number of hardware accelerated client-side received sent for a selected filter (entity).

hw-accel-client-bytes-in

The total number of hardware accelerated client-side received bytes for a selected filter (entity).

hw-accel-client-bytes-out

The total number of hardware accelerated client-side sent bytes for a selected filter (entity).

hw-accel-client-max-connections

The max number of hardware accelerated client-side connections for a selected filter (entity).

hw-accel-client-new-connections

The total number of hardware accelerated client-side new connections for a selected filter (entity).

hw-accel-client-current-connections

The average number of hardware accelerated client-side opened connections for a selected filter (entity).

hw-accel-server-packets-in

The total number of hardware accelerated server-side received packets for a selected filter (entity).

hw-accel-server-packets-out

The total number of hardware accelerated server-side received sent for a selected filter (entity).

hw-accel-server-bytes-in

The total number of hardware accelerated server-side received bytes for a selected filter (entity).

hw-accel-server-bytes-out

The total number of hardware accelerated server-side sent bytes for a selected filter (entity).

hw-accel-server-max-connections

The max number of hardware accelerated server-side connections for a selected filter (entity).

hw-accel-server-new-connections

The total number of hardware accelerated server-side new connections for a selected filter (entity).

hw-accel-server-current-connections

The average number of hardware accelerated server-side opened connections for a selected filter (entity).

tmm-total-ram

The average of total TMM memory for the selected filter (entity).

tmm-used-ram

The average of used TMM memory for the selected filter (entity).

tmm-free-ram

The average of free TMM memory for the selected filter (entity).

other-total-ram

The average of total other (daemons and kernel) memory for the selected filter (entity).

other-used-ram

The average of used other (daemons and kernel) memory for the selected filter (entity).

other-free-ram

The average of free other (daemons and kernel) memory for the selected filter (entity).

system-total-ram

The average of total system memory for the selected filter (entity).

system-used-ram

The average of used system memory for the selected filter (entity).

system-free-ram

The average of free system memory for the selected filter (entity).

swap-total-ram

The average of total swap memory for the selected filter (entity).

swap-used-ram

The average of used swap memory for the selected filter (entity).

swap-free-ram

The average of free swap memory for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override

Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-07-22 analytics vcmp report(1)

analytics vcmp scheduled-report

NAME

scheduled-report - Configure scheduled reports for vCMP.

MODULE

analytics vcmp

SYNTAX

Configure the scheduled-report component within the analytics vcmp module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the vCMP module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

```
list scheduled-report
```

Displays all of the vCMP scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics vcmp report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics vcmp report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics vcmp report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics vcmp scheduled-report(1)

analytics virtual report

NAME

report - Displays an virtual analytics report.

MODULE

analytics virtual

SYNTAX

Show, save or send an analytics virtual report using the syntax shown in the following sections.

DISPLAY

show report view-by [virtual]

options:

drilldown {

{

entity [virtual]

values

{

[value ...]

}

} ...

}

field-fmt

include-total

include-others

limit [number of rows]

measures {

[measure name ...]

}

order-by {

{

measure [measure name]

sort-type [asc / desc]

} ...

}

range [date range]

SAVE

save report view-by [virtual]

options:

drilldown {

```

{
entity [ virtual ]
values
{
[value ...]
}
} ...
}
file [ file name ]
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]

SEND
send-mail report view-by [ virtual ]
options:
drilldown {
{
entity [ virtual ]
values
{
[value ...]
}
} ...
}
email-addresses {
[email address ...]
}
format [ csv-aggregated | csv-time-series | pdf ]
include-total
include-others
limit [number of rows]
measures {
[measure name ...]
}
order-by {
{
measure [ measure name ]
sort-type [ asc / desc ]
} ...
}
}
range [date range]
smtp-config-override [ smtp configuration object name ]

```

DESCRIPTION

Use this command to generate virtual analytics reports. You can generate a VIP analytics report for the following entities:

• virtual - Virtual Server

Different measures are collected for each of these entities and can be a part of the report request.

EXAMPLES

```
show analytics virtual report view-by virtual
```

```
show analytics virtual report view-by virtual drilldown { { entity virtual values { 172.12.34.67 } } }
```

```
send-mail analytics virtual report view-by virtual measures { total-server-packets-in } limit 20 order-by { {
measure total-server-packets-in sort-type desc } } format pdf email-addresses { some.one@someaddress.com }
```

For more syntactical examples, see the tmsh help manual for analytics report.

OPTIONS

device

Specifies a BIG-IP device on which to generate a report. (Enterprise Manager only)

device-list

Specifies a custom list of BIG-IP devices on which to generate a report. (Enterprise Manager only)

drilldown

Specifies specific entities that are used as a filter.

email-addresses

Specifies the list of email addresses to which the report file is sent when using the send-mail command.

file Specifies the exported file path to be saved when using the save command. The file name should be simple

(not a full path).

format

Specifies the exported file format to be saved or sent. This option must be specified when using the save or send-mail commands.

include-others

Specifies that the grand total for the measure is displayed for all entities, except for those shown in the result. It can be used along with include-total.

include-total

Specifies that a total summary row should be added to the analytics report. For average measures, the total value is also an average.

limit

Specifies the maximum number of rows/entities in the output result set/file. The default value is 10, not including the total row/entity. The maximum value is 1000.

measures

Specifies a list of measures that can be used with the chosen entity type. The options are:

total-client-packets-in

The total number of client-side received packets for the selected filter (entity).

total-client-packets-out

The total number of client-side sent packets for the selected filter (entity).

total-server-packets-in

The total number of server-side received packets for the selected filter (entity).

total-server-packets-out

The total number of server-side sent packets for the selected filter (entity).

total-client-bits-in

The total number of client-side received bits for the selected filter (entity).

total-client-bits-out

The total number of client-side sent bits for the selected filter (entity).

total-server-bits-in

The total number of server-side received bits for the selected filter (entity).

total-server-bits-out

The total number of server-side sent bits for the selected filter (entity).

total-client-conns

The total number of client-side connections for the selected filter (entity).

avg-client-concurrent-conns

The average number of client-side concurrent connections for the selected filter (entity).

max-client-concurrent-conns

The max value of client-side concurrent connections for the selected filter (entity).

total-server-conns

The total number of server-side connections for the selected filter (entity).

avg-server-concurrent-conns

The average number of server-side concurrent connections for the selected filter (entity).

max-server-concurrent-conns

The max value of server side concurrent connections for the selected filter (entity).

total-syncookies

The total number of syncookies for the selected filter (entity).

total-syncookies-accepts

The total number of accepted syncookies for the selected filter (entity).

total-syncookies-rejects

The total number of rejected syncookies for the selected filter (entity).

total-hw-syncookies

The total number of HW syncookies for the selected filter (entity).

total-hw-syncookies-accepts

The total number of accepted HW syncookies for the selected filter (entity).

order-by

Specifies the measures and sort type (ascending or descending) that will be used to sort the final report. The value for each measure is a previously chosen measure. The default value for sort type is desc (descending).

range

Specifies the time/date range of the analytics information that you want to display. The given results will reflect the time range chosen here. The default value is the last hour (now--now-1h).

smtp-config-override
Specifies the SMTP configuration to use when sending reports by email. This overrides the default SMTP settings.

SEE ALSO

show, save, send-mail, tmsh, ltm profile analytics, analytics, analytics report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-07-23 analytics virtual report(1)

analytics virtual scheduled-report

NAME

scheduled-report - Configure scheduled reports for virtual.

MODULE

analytics virtual

SYNTAX

Configure the scheduled-report component within the analytics virtual module using the syntax shown in the following sections.

CREATE/MODIFY

```
create scheduled-report [name]
modify scheduled-report [name]
options:
  email-addresses [none | add | delete | modify |
    replace-all-with] { email-address [string] }
  first-time [date]
  frequency [every-6-hours | every-12-hours | every-24-hours | every-week | every-month]
  include-total [enabled | disabled]
  multi-leveled-report {
  chart-path [none | add | delete | modify | replace-all-with] { entity name [string] }
  limit [number of rows]
  time-diff [last-hour | last-day | last-week | last-month | last-year]
  view-by { entity name [string] }
  measures [none | add | delete | modify | replace-all-with] { measure name [string] }
  }
  predefined-report-name [name]
  smtp-config [name]
  device-group [name]
```

DISPLAY

```
list scheduled-report
list scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
show running-config scheduled-report
show running-config scheduled-report [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete scheduled-report [name]
```

DESCRIPTION

Use the scheduled-report component to create, modify or delete scheduled reports for the virtual module.

EXAMPLES

```
create scheduled-report myScheduledReport first-time now predefined-report-name "Top blocked URLs" frequency
every-6-hours email-addresses add { person@domain.com } smtp-config asm_smtp_conf
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top blocked URLs and is sent to person@domain.com using the smtp configuration defined in asm_smtp_conf.

```
modify scheduled-report myScheduledReport smtp-config none
```

Set smtp configuration of the scheduled report "myScheduledReport" to none, thus effectively disabling the scheduled report from begin generated and sent over eMail.

```
create scheduled-report myCustomScheduledReport first-time now email-addresses add { person@domain.com }
frequency every-6-hours smtp-config asm_smtp_conf multi-leveled-report { view-by url time-diff last-hour limit
5 chart-path add { policy violation } }
```

Creates a scheduled report, starting from the next hour and executing every 6 hours. The report contains a PDF showing statistics for the top 5 violated URLs after drilling-down to the top policy followed by the top

violation. The report is sent to person@domain.co using the smtp configuration defined in asm_smtp_conf.

list scheduled-report

Displays all of the virtual scheduled reports.

OPTIONS

email-addresses

A list of the email addresses of the recipients that receive the scheduled report.

first-time

First scheduled report time. Must be after current time and rounded up to the next round hour.

frequency

The scheduled report frequency. Example: every-6-hours means that the report will be generated and sent every 6 hours.

include-total

Enables or disables including a summary (Overall result) entity in results.

multi-leveled-report

Defines a custom multi-leveled report. Mutually exclusive with predefined-report-name. The multi-leveled-report definition contains the following parameters:

chart-path

A list of entities that define the scope in which the report will be displayed. For example: a chart path { violation url } means: Use the top violation list and generate a top URL list from it. These top URLs will be then used to display the view-by entity. For a list of valid entities see the help manual for analytics virtual report.

limit

The number of view-by entities displayed in the scheduled report.

time-diff

The time range for the report.

view-by

The main entity that the report is viewed by. For a list of valid entities see the help manual for analytics virtual report.

measures

The measures which are available for the selected entities.

predefined-report-name

Defines which predefined report (AKA predefined filter) will be used to generate the report. This keyword is mutually exclusive with multi-leveled-report.

smtp-config

Defines which SMTP configuration will be used to send the scheduled report. If set to none, the scheduled report will be disabled.

device-group

Defines the device-group which the report should generate the report for. If 'none' is set to this field, then the report will be generate for the 'self' device.

SEE ALSO

list, modify, show, tmsh, analytics virtual report, sys smtp-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2018-10-20 analytics virtual scheduled-report(1)

api-protection

api-protection profile apiprotection

NAME

apiprotection - Configures an API protection profile.

MODULE

api-protection profile

SYNTAX

Configure the apiprotection component within the profile module using the syntax shown in the following

sections.

CREATE/MODIFY

```
create apiprotection [name]
modify apiprotection [name]
options:
  access-profile [profile-access-name]
  app-service [[string] | none]
  default-response [response-name]
  default-server [[server-name] | none]
  defaults-from [apiprotection | [name]]
  description [[string] | none]
  dns-mode [ipv4-only | [ipv6-only] | [ipv6-prefer]]
  dns-resolver [[dns-resolver-name] | none]
  last-generated-path-id [integer]
  max-concurrent-subsessions [integer]
  openapi-version [[string] | none]
  paths [add | delete | none | replace-all-with] {
[path-name] {
  active [true | false]
  app-service [[string] | none]
  description [[string] | none]
  method [string]
  path-id [integer]
  server [[server-name] | none]
  uri [string]
}
}
  per-request-policy [per-request-policy-name]
  responses [add | delete | none | replace-all-with] {
[response-name]
}
  servers [add | delete | none | replace-all-with] {
[server-name]
}
  use-pool [false | true]
edit apiprotection [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
```

DISPLAY

```
list apiprotection
list apiprotection [ [ [name] | [glob] | [regex] ] ... ]
show running-config apiprotection
show running-config apiprotection [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  non-default-properties
  one-line
  recursive
```

```
show apiprotection
show apiprotection [name]
options:
  all
  default
  exa
  gig
  kil
  meg
  peta
  raw
  tera
  yotta
  zetta
```

DELETE

```
delete apiprotection [name]
options:
  all
  recursive
```

DESCRIPTION

You use the apiprotection component to configure an apiprotection profile. An API protection profile specified a group of settings that you can use to configure an API protection server.

NOTE: For the API protection profile to take effect, it must be associated with a virtual server that also specifies an HTTP profile.

EXAMPLES

```
create apiprotection myAPIProtectionProfile {
  access-profile myAPIProtectionProfile_ap
  default-response myAPIProtectionProfile_response1
  default-server myAPIProtectionProfile_server1
  defaults-from apiprotection
  description "My API protection Profile"
```

```

dns-mode ipv4-only
dns-resolver default-dns-resolver
last-generated-path-id 1
max-concurrent-subsessions 1
openapi-version "2.0"
partition Common
paths {
  myAPIProtectionProfile_path1 {
active true
method GET
path-id 1
uri /somepath
  }
}
per-request-policy myAPIProtectionProfile_prp
responses {
  myAPIProtectionProfile_response1
}
servers {
  myAPIProtectionProfile_server1
}
use-pool false
}

```

Creates an API protection profile named myAPIProtectionProfile based on the default profile named apiprotection. The profile provides protection to API requests handled by the server myAPIProtectionProfile_server1. Based on the Per-request-Policy configured in myAPIProtectionProfile_prp, a default response configured as myAPIProtectionProfile_response1 is provided for invalid requests. This profile serves GET requests to URI /somepath on the virtual server to which this profile is attached. The connection to the API server myAPIProtectionProfile_server1 is determined using the DNS resolver configuration default-dns-resolver resolving only IPv4 requests.

list apiprotection all all-properties

Displays a list of API protection profiles, including parameter values.

delete apiprotection myAPIProtectionProfile

Deletes the API protection profile named myAPIProtectionProfile.

OPTIONS

access-profile

Specifies the name of the associated access profile. If the API protection profile is created using REST API or GUI, the default access profile is automatically created and associated. The default is none if created using TMSH.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

default-response

Specifies the response name available under api-protection response. This value is mandatory. If the API protection profile is created using REST API or GUI, the default response is automatically created and associated using OpenAPI spec configuration. Create a response under api-protection response and associate here when using TMSH.

default-server

Specifies the server name available under api-protection server. If the API protection profile is created using REST API or GUI, default server is automatically created and associated.

defaults-from

Specifies the default API protection profile from which this profile is created. The default is apiprotection.

description

Specifies the description of the profile.

dns-mode

Specifies the DNS mode to use when resolving API server FQDN. Allowed values are ipv4-only, ipv6-only, and ipv6-prefer. The default is ipv4-only.

dns-resolver

Specifies the DNS resolver name configured under net dns-resolver. This cannot be empty when API Server is configured.

last-generated-path-id

Specifies the maximum path-id value configured for a path under paths. This value is used and set internally and requires no manual configuration.

max-concurrent-subsessions

Specifies the maximum number of concurrent subsessions. The default is 0, which sets the maximum number of concurrent subsessions to 5 times the licensed access session limit.

openapi-version

Specifies version information of the OpenAPI spec file used when creating the profile using REST API or GUI. This is set automatically when you use the spec file.

paths

Specifies the list of path configurations.

path-name

Specifies the name of the path configuration.

active

Specifies if the path-name is active. If path is inactive, Request-Classification-Agent under per-request policy will ignore the branch. The default is true.

description

Specifies description of path-name.

method

Specifies the HTTP method associated with the specific path path-name. This is mandatory input.

path-id

Specifies the path-id associated with the specific path path-name. This value is used in the Request Classification Agent under per-request policy to create a path specific branch.

server

Specifies the API server associated with the specific path path-name.

uri Specifies the URI associated with the specific path path-name. This is mandatory input.

per-request-policy

Specifies the per-request access policy attached to the API protection profile.

responses

Specifies the API response(s) associated with the profile. The configuration is defined under api-protection response.

servers

Specifies the API server name(s) associated with the profile. The configuration is defined under api-protection server.

use-pool

Specifies that the API protection profile is used to protect pool members.

SEE ALSO

api-protection response, api-protection server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2018. All rights reserved.

BIG-IP 2018-10-20 api-protection profile apiprotection(1)

api-protection response

NAME

response - Manages responses for API Protection Profile.

MODULE

api-protection response

SYNTAX

Configure the response component within the api-protection module using the following syntax.

CREATE/MODIFY

create response [name]

modify response [name]

options

app-service [[string] | none]

body [[string] | none]

description [[string] | none]

headers [add | delete | modify | none | replace-all-with] {

[name] {

app-service [[string] | none]

header-name [[string] | none]

header-value [[string] | none]

}

status-code [string]

status-string [string]

edit response [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line

DISPLAY

list response

list response [[[name] | [glob] | [regex]] ...]

show running-config response

show running-config response [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line

DELETE

delete response [name]

DESCRIPTION

You can use the response component to create and manage responses that can be used by the API Protection profile.

EXAMPLES

```
create response response401 {
description "Unauthorized Request"
status-code 401
status-string "Unauthorized"
}
```

Creates a response named response401 with status-code set to 401 and corresponding status-string set to Unauthorized.

```
list response
```

Displays a list of responses.

```
delete response response401
```

Deletes the response named response401.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

[name]

Specifies the name of the Response object. This setting is required.

body Specifies the body of the response. This value can be any string or session variable.

description

Specifies the description of the response.

partition

Displays the partition within which the component resides.

headers

Adds, deletes, modifies or replaces a set of headers, by specifying a header name and value for each entry. Header name and header value can be any string or session variable.

status-code

Specifies the response code that must be issued. This value can be any string or session variable. This setting is required.

status-string

Specifies the status string that must be issued. This value can be any string or session variable. This setting is required.

SEE ALSO

api-protection profile apiprotection

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-07-10 api-protection response(1)

api-protection server

NAME

server - Manages servers for API Protection Profile.

MODULE

api-protection server

SYNTAX

Configure the server component within the api-protection module using the following syntax.

CREATE/MODIFY

create server [name]

modify server [name]

options

app-service [[string] | none]

description [[string] | none]

serverssl-profile [[string] | none]

url [string]

edit server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DISPLAY

list server

list server [[[name] | [glob] | [regex]] ...]

show running-config server

show running-config server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete server [name]

DESCRIPTION

You can use the server component to create and manage servers that can be used by the API Protection profile.

EXAMPLES

```
create server serverA {
description "Server A"
serverssl-profile apm-default-serverssl
url "https://abc.com"
}
```

Creates a server named serverA with url https://abc.com and references a serverssl-profile.

```
list server
```

Displays a list of servers.

```
delete server serverA
```

Deletes the server named serverA.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

[name]

Specifies the name of the server object. This setting is required.

description

Specifies the description of the server.

partition

Displays the partition within which the component resides.

serverssl-profile

References a serverssl-profile. This setting is required when url is of https scheme.

url Specifies the URL of the server.

SEE ALSO

api-protection profile apiprotection

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-07-09 api-protection server(1)

apm

apm aaa active-directory-trusted-domains

NAME

active-directory-trusted-domains - Manages authentication access policy (AAA) Active Directory(r) Trusted Domains.

MODULE

apm aaa

SYNTAX

Configure the active-directory-trusted-domains component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create active-directory-trusted-domains [name]

modify active-directory-trusted-domains [name]

options:

app-service [[string] | none]

description [[string] | none]

root-domain [string]

trusted-domains [add | delete | modify | replace-all-with] {

{

active-directory [name]

}

}

edit active-directory-trusted-domains [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list active-directory-trusted-domains

list active-directory-trusted-domains [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete active-directory-trusted-domains [name]

DESCRIPTION

You can use the active-directory-trusted-domains component to manage AAA Active Directory Trusted Domains. You can use this object to configure cross-domain authentication across a forest. It also allows to configure Active Directory(r) agents to work in a Route Domains environment.

EXAMPLES

```
create active-directory-trusted-domains MyTRD { trusted-domains { myDomain1 myDomain2 myDomain3 } root-domain /Common/myDomain2 }
```

Creates an object named MyTRD, sets domains myDomain1, myDomain2, myDomain3 as trusted and the root-domain is set to myDomain2. To use this example you need to have Active Directory servers myDomain1, myDomain2 and myDomain3 pre-configured.

```
delete active-directory MyTRD
```

Deletes the AAA Active Directory Trusted Domains named MyTRD from the system.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

Specifies a user-defined description for the Active Directory Trusted Domains.

root-domain

Specifies an entry point to an Active Directory forest. An initial authentication request will always be sent to root domain first. This setting is required.

trusted-domains

Specifies a list of AAA Active Directory server components. Trust relationships should be defined for domains you add into this list. This setting is required.

SEE ALSO

active-directory

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2016-01-07 apm aaa active-directory-trusted-domains(1)

apm aaa active-directory

NAME

active-directory - Manages an authentication access policy (AAA) Active Directory(r) server.

MODULE

apm aaa

SYNTAX

Configure the active-directory component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create active-directory [name]

modify active-directory [name]

options:

admin-encrypted-password [[string] | none]

admin-name [[string] | none]

app-service [[string] | none]

cleanup-cache [pso | group | kerberos | none]

description [[string] | none]

domain [[string] | none]

domain-controller [[string] | none]

domain-controllers [add | delete | modify | replace-all-with] {

[name] {

ip [ip address]

}

}

group-cache-ttl [integer]

domain-controllers none

location-specific [true | false]

pool [name]

pso-cache-ttl [integer]

padata [encryption type]

timeout [integer]

edit active-directory [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list active-directory

list active-directory [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete active-directory [name]

DESCRIPTION

You can use the active-directory component to manage an AAA Active Directory server. The Active Directory is a network structure supported by Windows(r) 2000, or later, that provides support for tracking and locating any object on a network.

EXAMPLES

```
create active-directory MyADserver { domain-controller "server01.company.com domain "company.com " admin-name "administrator" admin-encrypted-password "!My123Password" }
```

Creates the AAA Active Directory server named MyADserver in the company.dom domain, sets the administrator logon name to administrator and the administrator password to !My123Password, and sets the Key Distribution Center to company.com.

```
delete active-directory MyActiveDirectoryServer
```

Deletes the AAA Active Directory server named MyActiveDirectoryServer from the system.

OPTIONS

`admin-encrypted-password`

Specifies the password associated with admin name. This option is required only when you are using an Active Directory Query agent with this Active Directory server object.

`admin-name`

Specifies the user name that has administrative permissions on an AAA Active Directory server. This option is required only when you are using an Active Directory Query agent with this Active Directory server object.

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`cleanup-cache`

Specifies whether this is a cache cleanup request. You can clean up the group, PSO or Kerberos cache. The default value is none.

`description`

Specifies a description for the component. The default is none.

`domain`

Specifies the Active Directory domain name. This setting is required.

`[name]`

Specifies the name of an AAA Active Directory server. This setting is required.

`domain-controller`

Specifies the fully qualified domain name (FQDN) of the domain controller for the domain specified in the domain option. The default is none.

`domain-controllers`

Adds, deletes, or replaces a set of domain controllers, by specifying an FQDN for each entry. You can configure the following options for each domain controller:

`ip` An IP address for specified domain controller entry.

`group-cache-ttl`

Specifies group cache lifetime in days [0..1825]. The default value is 30. If you specify group cache lifetime 0, that means cache will be updated on every request.

`pso-cache-ttl`

Specifies password security objects (PSO) Cache lifetime in days [0..1825]. The default value is 30. If you specify PSO cache lifetime 0, that means cache will be updated on every request.

`location-specific`

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

`pool` Specifies the name of the pool with which the server is associated. The default is none.

`partition`

Displays the partition within which the component resides. The default is Common.

`padata`

Specifies a Kerberos preauthentication encryption type. If it is specified, the BIG-IP system includes Kerberos preauthentication data within the first AS-REQ. If you do not need to include preauthentication data, set this option to "none". Supported encryption types: none, des-cbc-crc, des-cbc-md5, aes128-cts-hmac-sha1-96, aes256-cts-hmac-sha1-96, rc4-hmac. The default is rc4-hmac.

`timeout`

Specifies a timeout interval (in seconds) after which an AAA Active Directory server closes a connection. The default is 15.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

apm aaa crldp

NAME

crldp - Configure a Certificate Revocation List Distribution Point (CRDLP) server object for implementing a CRLDP authentication module.

MODULE

apm aaa

SYNTAX

Configure the crldp component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create crldp [name]

modify crldp [name]

options:

address [ip addr]

allow-nullcrl [true | false]

app-service [[string] | none]

base-dn [[string> | none]

cache-expire [[integer] | none]

connection-timeout [[integer] | none]

description [[string> | none]

location-specific [true | false]

pool [name]

port [[integer] | none]

reverse-dn [true | false]

use-issuer [true | false]

use-pool [enabled | disabled]

verify-sig [true | false]

edit crldp [[glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list crldp

list crldp [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete crldp [name]

DESCRIPTION

Configure a CRLDP authentication server, and then assign the server to the CRLDP auth agent in your access policy.

EXAMPLES

```
create crldp aaa-ldap-2027 { address 172.27.32.60 allow-nullcrl false base-dn DC=net,DC=aina,DC=test cache-
expire 1000 connection-timeout 15 description none partition Common pool aaa-ldap-2027-pool port ldap reverse-
dn true use-issuer false use-pool disabled verify-sig true }
```

Creates a CRLDP server named aaa-ldap-2027.

```
delete crldp server my_crldp_server
```

Deletes the CRLDP server named my_crldp_server.

OPTIONS

address

Specifies the IP address of the server. This option is required.

allow-nullcrl

Specifies whether to consider a null CRL from the CRLDP server a successful authentication. The default is false.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

base-dn

Specifies the LDAP base directory name for certificates that specify the CRL distribution point in directory name (dirName) format. Used when the value of the X509v3 attribute crlDistributionPoints is of type dirName. In this case, the BIG-IP system attempts to match the value of the crlDistributionPoints

attribute to the Base DN value. An example of a Base DN value is cn=lxxx,dc=f5,dc=com.

cache-expire

Specifies (in seconds) an update interval for CRL distribution points. The update interval for distribution points ensures that CRL status is checked at regular intervals, regardless of the CRL timeout value. This helps prevent CRL information from becoming outdated before the Access Policy Manager checks the status of a certificate.

connection-timeout

Specifies the number of seconds of inactivity the system allows before the connection times out. The default is 15.

description

Specifies a unique description for the server. The default is none.

partition

Displays the partition within which the component resides.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

pool Specifies the name of the pool with which the server is associated.

port Specifies the CRLDP service port. The default is 389.

reverse-dn

Specifies in which order the system is to attempt to match the Base DN value to the value of the X509v3 attribute crlDistributionPoints. Possible values are enabled and disabled. When set to enabled, the system matches the base DN from left to right, or from the beginning of the DN string, to accommodate dirName strings in certificates such as C=US,ST=WA,L=SEA,OU=F5,CN=xxx. The default value is false.

use-issuer

Specifies whether the CRL distribution point is extracted from the certificate of the client certificate issuer. The default is false.

use-pool

Enables or disables high availability between CRLDP servers. When enabled, Access Policy Manager sends CRLDP authentication requests for the associated CRLDP auth agent to the virtual server, and standard pool behavior is used to implement high availability for CRDLP.

verify-sig

Specifies whether the signature on the received CRL is verified. The default if true.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 apm aaa crldp(1)

apm aaa endpoint-management-system

NAME

endpoint-management-system - Manages an integration with a remote Mobile Device Management (MDM) server.

MODULE

apm aaa

SYNTAX

Configure the endpoint-management-system component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create endpoint-management-system [name]

options:

type [airwatch | maas360 | ms-intune]

fqdn [string]

port [port]

serverssl-profile [name]

description [[string] | none]

username [string]

password [string]

mdm-token [[string] | none]

billing-id [[string] | none]
application-id [[string] | none]
access-key [[string] | none]
platform [[string] | none]
tenant-id [[string] | none]
client-id [[string] | none]
client-secret [[string] | none]
dns-resolver [[name] | none]
app-version [[string] | none]
sync-interval [[integer] | none]
location-specific [true | false]
modify endpoint-management-system [name]

options:

fqdn [string]
port [port]
serverssl-profile [name]
description [[string] | none]
username [string]
password [string]
mdm-token [[string] | none]
billing-id [[string] | none]
application-id [[string] | none]
access-key [[string] | none]
platform [[string] | none]
tenant-id [[string] | none]
client-id [[string] | none]
client-secret [[string] | none]
dns-resolver [[name] | none]
app-version [[string] | none]
sync-interval [[integer] | none]
location-specific [true | false]

edit endpoint-management-system [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list endpoint-management-system

list endpoint-management-system [[[name] | [glob] | [regex]] ...]

options:

all-properties
app-service
non-default-properties
one-line
partition

DELETE

delete endpoint-management-system [name]

DESCRIPTION

You can use the endpoint-management-system component to manage an integration with remote Mobile Device Management (MDM) server.

EXAMPLES

```
create endpoint-management-system MyEndpointManagementSystem { type airwatch fqdn "server01.company.com" port 443 serverssl-profile serverssl username "administrator" password "!My123Password" mdm-token "token" }
```

Creates the endpoint management system named MyEndpointManagementSystem with MDM API URL server01.company.com:port, sets serverssl-profile to serverssl, sets the MDM administrator user name to administrator and the password to !My123Password, and sets the API token to token.

```
delete endpoint-management-system MyEndpointManagementSystem
```

Deletes the endpoint management system named MyEndpointManagementSystem from the system.

OPTIONS

name Specifies the name for the endpoint management system. This setting is required.

type Specifies the type of endpoint management system: airwatch, maas360 or ms-intune. This setting is required.

fqdn Specifies the fully qualified domain name. This setting is required.

port Specifies the port number. Default is 443

serverssl-profile
Specifies the server SSL profile. This setting is required.

description
Specifies a description for the component. The default is none.

username
Specifies the user name of the MDM administrator. This setting is required.

password
Specifies the password the MDM administrator uses to log in. This setting is required.

mdm-token

Specifies the API token.

billing-id

Specifies the billing ID for the user's IBM Maas360 account. Valid only for maas360 type.

application-id

Specifies the application ID provided by IBM Maas360. Valid only for maas360 type.

access-key

Specifies the access key provided by IBM Maas360. Valid only for maas360 type.

platform

Specifies the platform version of the IBM Maas360 console. Valid only for maas360 type.

tenant-id

Specifies the Microsoft Intune Tenant Id. Valid only for ms-intune type.

client-id

Specifies the Client Id of the Web App created on Microsoft Azure for the Microsoft Intune Integration. Valid only for ms-intune type.

client-secret

Specifies the Client Secret of Web App created on Microsoft Azure for the Microsoft Intune Integration. Valid only for ms-intune type.

dns-resolver

Specifies the Dns Resolver. Valid only for ms-intune type.

app-version

Specifies the current version number of the application that corresponds to the account.

sync-interval

Specifies the length of time it takes for the synchronization to complete. The default is 240 minutes.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015. All rights reserved.

BIG-IP 2018-01-18 apm aaa endpoint-management-system(1)

apm aaa f5-mfa-configuration

NAME

f5-mfa-configuration - defines F5 multi-factor authentication configuration.

MODULE

apm aaa

SYNTAX

Configure the f5-mfa-configuration component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

```
create f5-mfa-configuration [name]
modify f5-mfa-configuration [name]
options:
  app-service [[string] | none]
  f5-service-connector [name]
  permitted-devices-types [add | delete | modify | replace-all-with] {
    [mobile | totp]
  }
  max-mobile-devices-per-user [[integer] | none]
  registration-sms-template [[string] | none]
  require-biometric [[true | false] | none]
```

```
edit f5-mfa-configuration [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list f5-mfa-configuration
list f5-mfa-configuration [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete f5-mfa-configuration [name]
```

DESCRIPTION

You can use the f5-mfa-configuration component to define F5 multi-factor authentication configuration.

EXAMPLES

```
create f5-mfa-configuration MyF5MFAConfiguration { f5-service-connector MyF5ServiceConnector permitted-
devices-types { mobile } max-mobile-devices-per-user 2 registration-sms-template "Hello, Please follow the
link below to register your device for second factor authentication:
```

```
  %{session.f5_mfa.device_registration.registration_url}" require-biometric true }
Creates the f5_mfa configuration named MyF5MFAConfiguration with f5-service-connector
MyF5ServiceConnector, adds mobile to permitted-devices-types, sets max-mobile-devices-per-user to 2, sets
registration-sms-template to Hello, Please follow the link below to register your device for second
factor authentication: %{session.f5_mfa.device_registration.registration_url} and sets require-biometric
to true
```

```
delete f5-mfa-configuration MyF5MFAConfiguration
Deletes the f5_mfa configuration named MyF5MFAConfiguration from the system.
```

OPTIONS

```
[name]
Specifies the name for the f5 mfa configuration. This setting is required.
```

```
f5-service-connector
Specifies the f5-service-connector. This setting is required.
```

```
permitted-devices-types
Specifies permission of the use of mobile devices or hardware tokens (TOTP) or both for multi-factor
authentication. This setting is required.
```

```
max-mobile-devices-per-user
Specifies the number of devices that one user can register for multi-factor authentication.
```

```
registration-sms-template
Specifies the message to send to a user to register their mobile devices.
```

```
require-biometric
Set this item to true to require that the user present a physical characteristic, such as a fingerprint,
on the mobile device for an additional authentication factor.
```

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2017. All rights reserved.

BIG-IP 2017-09-19 apm aaa f5-mfa-configuration(1)

apm aaa f5-service-connector

NAME

f5-service-connector - Specifies properties for establishing a connection to an F5 multi-factor authentication service that is external to and separate from BIG-IP Access Policy Manager and the BIG-IP system.

MODULE

apm aaa

SYNTAX

Configure the f5-service-connector component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

```
create f5-service-connector [name]
modify f5-service-connector [name]
options:
```

app-service [[string] | none]
service-url [string]
customer-id [string]
customer-key [string]
dns-resolver [name]
serverssl-profile [name]

edit f5-service-connector [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties

DISPLAY
list f5-service-connector
list f5-service-connector [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 app-service
 non-default-properties
 one-line
 partition

DELETE
delete f5-service-connector [name]

DESCRIPTION
 You can use the f5-service-connector component to create F5 MFA Configuration component.

EXAMPLES
 create f5-service-connector MyF5ServiceConnector { service-url "https://service01.company.com" customer-id "id01" customer-key "key01" dns-resolver new-dns-resolver serverssl-profile serverssl }
 Creates the F5 service connector named MyF5ServiceConnector with service URL https://service01.company.com, sets the customer id to id01, sets the customer-key to key01, sets dns-resolver to new-dns-resolver and sets serverssl-profile to serverssl

 delete f5-service-connector MyF5ServiceConnector
 Deletes the F5 service connector named MyF5ServiceConnector from the system.

OPTIONS
 [name]
 Specifies the name for the F5 service connector. This setting is required.

 app-service
 Specifies the name of the application service to which the object belongs. The default value is none.
 Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

 service-url
 Specifies the URL for connecting to the service. This setting is required.

 customer-id
 Specifies the customer id for the service. This setting is required.

 customer-key
 Specifies the customer key for the service. This setting is required.

 dns-resolver
 Specifies the DNS resolver for the connector to use. This setting is required.

 serverssl-profile
 Specifies the server SSL profile. This setting is required.

SEE ALSO
COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2017. All rights reserved.

BIG-IP 2017-09-19 apm aaa f5-service-connector(1)

apm aaa http-connector-request

NAME
 http-connector-request - Stores the configuration for a HTTP Request.

MODULE
 apm aaa

SYNTAX

Configure the http-connector-request object within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

```
create/modify http-connector-request [name]
```

options:

method [string]

url [string]

transport [name]

auth [none | basic | bearer | custom]

username [[string] | none]

password [[string] | none]

token [[string] | none]

request-headers [[string] | none]

request-body [[string] | none]

response-action [parse | save | ignore]

response-headers [[string] | none]

```
edit http-connector-request [ [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties

non-default-properties

DISPLAY

```
list http-connector-request
```

```
list http-connector-request [ [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

```
delete http-connector-request [name]
```

DESCRIPTION

You can use the http-connector-request to send HTTP Request to any Endpoint.

EXAMPLES

```
create http-connector-request PostRequest { method POST url
```

```
"https://server01.company.com?id=%{subsession.mdm.device_id}" transport "your-http-connector-transport" auth
```

```
bearer token some-token response-action ignore }
```

Creates the HTTP Connector Request Object named PostRequest with HTTP method POST, sets the URL to https://server01.company.com?id=%{subsession.mdm.device_id}, sets transport to your-http-connector-transport, sets the Authentication Type to bearer and the token to some_token, and sets the Response Action to ignore.

```
delete http-connector-request PostRequest
```

Deletes the HTTP Connector Request Object named PostRequest from the system.

OPTIONS

name Specifies the name for the HTTP Connector Request Object. This setting is required.

method

Specifies the method for http request, e.g. GET, POST, PUT etc. This setting is required.

url Specifies the url of the endpoint. This setting is required. Sub-session variables can be specified,

%{subsession.variable_name}.

transport

Specifies HTTP Connector Transport object. This setting is required.

auth Specifies the Authentication type. Default is none. If custom is selected then username, password and token fields can be used in url, request-headers or request-body with standard placeholders for them e.g.

%{password}, %{username} or %{token}. Along with this special placeholder %{basic_auth} can be used to insert value "Basic base64_encoded_value[some-username:some-password]" where "some-username" and "some-password" is specified in username and password field.

username

Specifies the user name. This field can be used in url, request-headers or request-body with placeholder %{username} if Authentication type is custom.

password

Specifies the password. This field can be used in url, request-headers or request-body with placeholder %{password} if Authentication type is custom.

token

Specifies the token. This field can be used in url, request-headers or request-body with placeholder %{token} if Authentication type is custom.

request-headers

Specifies the request headers. Each header needs to be specified on a new-line. There is no easy way to specify new-line separating the headers if creating the object through tmsh. Since this is not a required field, one can create the http query without this field and then use edit command on the created object

to specify the headers each separated through new-line. Sub-session variables can be specified, `%{subsession.variable_name}`.

request-body
Specifies the request body. Sub-session variables can be specified, `%{subsession.variable_name}`.

response-action
Specifies how response data is processed. Action 'parse' should be selected only if response is in JSON format. 'save' will save the entire response data in a single sub-session variable. 'ignore' will not save the response.

response-headers
Specifies the response headers which needs to be saved as sub-session variables. These are "," separated values. If it is empty then response headers are ignored.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015. All rights reserved.

BIG-IP 2019-06-26 apm aaa http-connector-request(1)

apm aaa http-connector-transport

NAME

http-connector-transport - stores the Network level configuration used by HTTP Connector Request.

MODULE

apm aaa

Configure the http-connector-transport object within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create/modify http-connector-transport [name]

options:

ssl-profile [name]
dns-resolver [[name] | none]
max-response-size [integer]
timeout [integer]

edit http-connector-transport [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list http-connector-transport

list http-connector-transport [[[name] | [glob] | [regex]] ...]

options:

all-properties
app-service
non-default-properties
one-line
partition

DELETE

delete http-connector-transport [name]

DESCRIPTION

You can use the http-connector-transport to configure Network level configuration used by HTTP Connector Request.

EXAMPLES

```
create http-connector-transport HttpConnector { ssl-profile serverssl dns-resolver your-dns-resolver max-response-size 32768 timeout 5 }
```

Creates the HTTP Connector Transport Object named HttpConnector with Server SSI Profile serverssl, sets the DNS resolver your-dns-resolver, sets the Max Response Size to 32768 bytes and sets the Timeout to be 5 seconds.

```
delete http-connector-transport HttpConnector
```

Deletes the HTTP Connector Transport Object named HttpConnector from the system.

OPTIONS

name Specifies the name for the HTTP Connector Transport Object. This setting is required.

ssl-profile
Specifies the server SSL profile.

dns-resolver
Specifies the Dns Resolver. This setting is required.

max-response-size
Specifies the limit on size of the response body that is allowed in bytes. If response body size is greater than specified limit then the data in body is ignored. If response headers were specified in http-connector-request object, they will still be processed in the event of response body being ignored due to max-response-size off limit. Default is 32768 bytes.

timeout
Specifies the timeout in seconds.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015. All rights reserved.

BIG-IP 2019-06-26 apm aaa http-connector-transport(1)

apm aaa http

NAME

http - Specify an http server configuration used for authentication.

MODULE

apm aaa

SYNTAX

Configure the http component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create http [name]

modify http [name]

options:

app-service [[string] | none]

auth-type [form-based | basic-ntlm | custom-post]

content-type [xml-utf8 | url-encoded-utf8 | none]

custom-body [[string] | none]

description [[string] | none]

follow-redirect [integer]

form-action [[string] | none]

form-fields [[string] | none]

form-method [get | post]

form-params [[string] | none]

form-password [[string] | none]

form-username [[string] | none]

headers [add | delete | modify | replace-all-with | none] {

[name] {

app-service [[string] | none]

hname [[string] | none]

hvalue [[string] | none]

}

location-specific [true | false]

start-uri [[string] | none]

success-match-type [url | cookie | string | exact-cookie]

success-match-value [[string] | none]

edit http [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list http

list http [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE
delete http [name]

DESCRIPTION

You can use the http component to create and manage AAA HTTP servers.

EXAMPLES

```
create http myHttpServer { start-uri "http://mycompany.com/" auth-type basic-ntlm }
```

Creates an HTTP authentication server named "myHttpServer" with a starting URI of http://mycompany.com.

```
delete http myHttpServer
```

Deletes the myHttpServer AAA HTTP server.

OPTIONS

- app-service**
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.
- auth-type**
Specifies the type of authentication you want to use.
- form-based**
Specifies the authentication type to be form-based.
- basic-ntlm**
Specifies the authentication type to be basic-ntlm.
- custom-post**
Specifies the authentication type to be custom-post.
- content-type**
Specifies the encoding (xml-utf8, url-encoded-utf8, or none) for an HTTP custom post. If you specify 'none', you must use the headers option to add a custom header. In addition to specifying a custom header, you must apply your own encoding through an iRule.
- custom-body**
Specifies the body for a HTTP Custom Post.
- description**
Specifies a unique description for the server. The default is none.
- follow-redirect**
Specifies the number of pages away from the landing page the request should travel before failing.
- form-action**
Specifies the complete destination URL to process the form using HTTP form-based authentication. This is optional. If you do not specify a form action, then Access Policy Manager will use the URI from the request to perform HTTP form-based authentication.
- form-fields**
Specifies the hidden form parameters that are required by the authentication server logon form at your location. The default is none. Specify a parameter name, a space, and the parameter value, if any. Multiple parameters can be configured with each "name value" pair in one line. Use edit to add multiple parameters. Please note that create and modify do not allow using new line on the terminal.
- form-method**
Specifies the form method you want to use for the form-based HTTP authentication. The value is either Get or POST. The default is POST. However, if you specify GET, the Access Policy Manager will force the authentication using HTTP GET rather than perform authentication using form-based POST.
- form-password**
Specifies the parameter names used by the form you are sending the POST request to.
- form-username**
Specifies the parameter names used by the form you are sending the POST request to.
- headers**
Specifies the name and value of the header content to be inserted in an HTTP Post. The options are:
 - app-service**
Specifies the name of the application service to which the HTTP header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the HTTP header. Only the application service can modify or delete the HTTP header.
 - hname**
The name of the HTTP header.
 - hvalue**
The value of the HTTP header.
 - location-specific**
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies the name of the aaa http server. This option is required.

partition

Displays the partition within which the component resides. The default is Common.

start-uri

Specifies a URL resource, for example, `http://plum.tree.lab2.sp.companynet.com/`. This resource must respond with a challenge to a non-authenticated request.

success-match-type

Specifies the method your authentication server uses and determines the option definition used for this field. The field toggles according to your selection.

cookie

Specifies any string in cookie is required.

exact-cookie

Specifies key fields in cookie is required.

string

Specifies a specific string is required.

url Specifies a URL is required.

success-match-value

Specifies the URL, any string in cookie, exact cookie or specific string used for the specific success match type you see.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm aaa http(1)

apm aaa kerberos-keytab-file

NAME

kerberos-keytab-file - Manages a Kerberos keytab file.

MODULE

apm aaa

SYNTAX

Configure the kerberos-keytab-file component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create kerberos-keytab-file [name]

modify kerberos-keytab-file [name]

options:

app-service [[string] | none]

source-path [string]

edit kerberos-keytab-file | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list kerberos-keytab-file

list kerberos-keytab-file [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete kerberos-keytab-file [name]

DESCRIPTION

You can use the kerberos-keytab-file component to create and manage a Kerberos Keytab file.

EXAMPLES

```
create kerberos-keytab-file my_keytab { source-path file:/root/apmkeytab }
Creates a Kerberos Keytab file name my_keytab located at root/apmkeytab.
```

```
delete kerberos-keytab-file my_keytab
Deletes the Kerberos Keytab file name my_keytab.
```

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

source-path
Specifies the location of the Kerberos Keytab file.

partition
Displays the partition within which the component resides.

SEE ALSO

apm aaa kerberos, apm policy agent kerberos

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2014-10-27 apm aaa kerberos-keytab-file(1)

apm aaa kerberos

NAME

kerberos - Configures a Kerberos server.

MODULE

apm aaa

SYNTAX

Configure the kerberos component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

```
create kerberos [name]
modify kerberos [name]
options
  auth-realm [[string] | none]
  app-service [[string] | none]
  keytab-file-obj [[string] | none]
  location-specific [true | false]
  service-name [[string] | none]
```

```
edit kerberos [ [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list kerberos
list kerberos [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete kerberos [name]
```

DESCRIPTION

You can use the kerberos component to create and manage AAA Kerberos servers. Use the Kerberos authentication server to configure authentication for the Access Policy Manager. A client retrieves credentials from the domain controller and passes those credentials to the Access Policy Manager. Then Access Policy Manager uses the value in the keytab-file-obj option of the Kerberos AAA server object to verify the credentials. Access Policy Manager system does not have to reside in the domain.

EXAMPLES

delete kerberos my_kerberos
Deletes the server named my_kerberos.

OPTIONS

auth-realm

Specifies a Kerberos auth realm name (administrative name), such as user@realm.com to establish the boundaries within which an authentication server has the authority to authenticate a user, host, or service. Kerberos clients manually map DNS domain names to Kerberos realm names. This option is required.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

keytab-file-obj

Specifies a keytab file that contains the keys (derived from the Kerberos password) that the server uses to authenticate the client. This option is required.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies the name of an AAA Kerberos server. This option is required.

partition

Displays the partition within which the component resides.

service-name

Specifies the Kerberos service name defined inside KDC in the format service name/hostname@kerberosrealm. This option is required, for example, HTTP.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2014-10-27 apm aaa kerberos(1)

apm aaa ldap

NAME

ldap - Manages an AAA LDAP server.

MODULE

apm aaa

SYNTAX

Configure the ldap component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create ldap [name]

modify ldap [name]

options:

address [[ip addr] | none]

admin-dn [[string] | none]

admin-encrypted-password [[string] | none]

app-service [[string] | none]

base-dn [string]

description [[string] | none]

is-ldaps [false | true]

location-specific [true | false]

pool [name]

port [[service] | none]

schema-attr {

group-member [[string] | none]

group-member-value [[string] | none]

group-memberof [[string] | none]

group-object-class [[string]]

user-memberof [[string] | none]

user-object-class [string]

}

serverssl-profile [none | serverssl | serverssl-insecure-compatible | wom-default-serverssl]

timeout [integer]

use-pool [enabled | disabled]

edit ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ldap

list ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete ldap [name]

DESCRIPTION

You can use the ldap component to create and manage an AAA LDAP server.

EXAMPLES

```
create ldap MyLDAPserver { address 172.30.6.144 admin-dn
```

```
"cn=administrator,cn=users,dc=company,dc=companynet,dc=com" admin-encrypted-password "!MyPassword" }
```

Creates the AAA LDAP server named MyLDAPserver that is assigned the IP address 172.30.6.144 and the cn=administrator,cn=users,dc=company,dc=companynet,dc=com admin dn with a password of !MyPassword.

```
delete ldap MyLDAPServer
```

Deletes the AAA LDAP server named MyLDAPServer from the system.

OPTIONS

address

Specifies the IP address of an AAA LDAP server. This option is required.

admin-dn

Specifies the Container Distinguished Name (DN) to use for authentication. This option is required.

admin-encrypted-password

Specifies the password for admin name. This option is required.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

base-dn

Specifies the base DN from which to search. This search DN is used to search groups across a whole directory.

group-cache-ttl

Specifies a lifetime for the group cache (days).

cleanup-cache

Specifies whether cache invalidation is required. The default is none. The options are:

none

group

description

Specifies a unique description for the server. The default is none.

is-ldaps

Specifies whether to use the LDAPS protocol during authentication. If true, you must also specify the option serverssl-profile.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies the name of the AAA server. This option is required.

partition

Displays the partition within which the component resides.

pool Specifies the name of the pool with which the server is associated. The default is none.

port Specifies the port number of the AAA LDAP server. The default is ldap. This option is required.

schema-attr

Specifies LDAP schema-specific attribute names.

user-object-class The value of the objectClass attribute for a user object. The default is "user".

user-memberof If the user object maintains a group membership, you should specify the membership

attribute name here. The default is "memberOf".
group-object-class The value of the objectClass attribute for a group object. The default is "group".
group-memberof If the group object maintains a group membership in other groups, you should specify a membership attribute name here. The default is "memberOf".
group-member If the group object maintains a list of users that belong to the group, you should specify the attribute here. The default is "member".
group-member-value If the "group-member" attribute is specified, you should specify the attribute that is used to add users into a group. The default is "dn".
serverssl-profile
Specifies the server side SSL profile. LDAPS is achieved by directing LDAP traffic over a virtual server that uses a server side SSL to communicate with the LDAP server.

The options are:

serverssl
serverssl-insecure-compatible
wom-default-serverssl
timeout
Specifies a timeout interval (in seconds) for the AAA server after which the server closes a connection. The default is 15.

use-pool
Enables or disables high availability between pool members. When enabled, the Access Policy Manager sends AAA requests for the associated policy item to the virtual server, and standard pool behavior is used to implement high availability for CRDLP.

SEE ALSO
COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2014-10-27 apm aaa ldap(1)

apm aaa oam

NAME
oam - Manages an AAA Oracle Access Manager server.

MODULE
apm aaa

SYNTAX
Configure the oam component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY
create oam [name]
modify oam [name]
options:
access-server-hostname [[string] | none]
access-server-name [[string] | none]
access-server-port [[integer] | none]
access-server-retries [integer]
accessgate-encrypted-password [[string] | none]
accessgates [add | delete | modify | replace-all-with] {
[name]
}
action [config-accessgate | noop]
admin-id [[string] | none]
admin-password [[string] | none]
app-service [[string] | none]
description [[string] | none]
enable [false | true]
global-access-protocol-passphrase [[string] | none]
location-specific [true | false]
transport-security-mode [cert | open | simple]

edit oam | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list oam
list oam [[name] | [glob] | [regex]] ...]
options:
all-properties

app-service
non-default-properties
one-line
partition

DELETE
delete oam [name]

DESCRIPTION
You can use the oam component to create and manage an AAA Oracle Access Manager server.

EXAMPLES
create oam oam10g { access-server-hostname www.localcorp.biz access-server-name accessSrv1 access-server-port 6021 access-server-retries 0 accessgates { oam10gwebgate1 { encrypted-password [string] } } admin-id firstname.lastname admin-password "[string]" global-access-protocol-passphrase "[string]" transport-security-mode simple }
Creates the AAA OAM server named oam10g accessing the web gate oam10gwebgate1 on the Access Server accessSrv1 at host name www.localcorp.biz on port 6021. The server retries connections zero times.

delete aaa oam MyOAMServer
Deletes the AAA Oracle Access Manager server named MyOAMServer from the system.

OPTIONS

access-server-hostname
Specifies the IP address or FQDN of the Oracle Access Manager server. This option is required.

access-server-name
Specifies the name of the Oracle Access Manager server. This option is required.

access-server-port
Specifies the port of the Oracle Access Manager server. The default is 6021.

access-server-retries
Specify the number of times you want the access gate to attempt to connect to the Oracle Access Manager server when the action option is set to config-accessgate. The default is 0 (zero).

accessgates
Specifies the ID of the access gate or web gate on the OAM Server. The system supports the use of multiple access gates/web gates as long as they are from the same OAM server.

action
Specifies the Oracle Access Manager action type. Actions allow you to pass user profile information or to redirect the user's browser to another site. For more information on Actions, refer to the Access Administration Guide provided by Oracle. The options are:

config-accessgate
Specifies that you want the system to use the configureAccessGate tool.

noop Specifies "no operation performed." This is the default.

admin-id
Specifies the administrator ID required by the Oracle Access Manager server. This option is required.

admin-password
Specifies the administrator password required by the Oracle Access Manager server. The default is none.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description
Specifies a unique description for the Oracle Access Manager server. The default is none.

enable
Specifies whether you want to enable the server. The default is true.

global-access-protocol-passphrase
Specifies a global passphrase for all Oracle components. The default is none.

location-specific
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]
Specifies the name of an AAA Oracle Access Manager server. This setting is required.

transport-security-mode
Specifies the transport security level for the communication between Oracle components and Access Policy Manager. The options are:

open Communication is not encrypted for protection. Use this mode when security is not an issue

simple
Communication is encrypted with Oracle Access Manager's internal CA. Simple mode encrypts

communications using Transport Layer Security, RFC 2246 (TLS v1). This mode is less secure than Cert mode. Use this mode if you have some security concerns but do not want to manage your own CA.

cert Communication is encrypted with an external CA. Use cert mode if you want different certificates on OAM servers and webgates and you have a trusted 3rd party CA. Oracle Access Manager components use X.509 digital certificates in PEM format only.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2014-10-27 apm aaa oam(1)

apm aaa oauth-provider

NAME

oauth-provider - Manages an OAuth Provider.

MODULE

apm aaa

SYNTAX

Configure the oauth-provider component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create oauth-provider [name]

modify oauth-provider [name]

options:

allow-self-signed-jwk-cert [bool]

app-service [[string] | none]

authentication-uri [[string] | none]

auto-jwt-config-name [[string] | none]

description [[string] | none]

ignore-expired-cert [bool]

last-discovery-time [date/time]

manual-jwt-config-name [[string] | none]

max-json-nesting-layers [integer]

max-response-size [integer]

openid-cfg-uri [[string] | none]

save-json-payload [enabled | disabled]

token-uri [[string] | none]

token-validation-scope-uri [[string] | none]

trusted-ca-bundle [[string] | none]

type [custom | f5 | facebook | google | ping]

use-auto-jwt-config [bool]

userinfo-request-uri [[string] | none]

edit oauth-provider [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list oauth-provider

list oauth-provider [[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete oauth-provider [name]

DESCRIPTION

You can use the oauth-provider component to manage an OAuth Provider. The OAuth Provider specifies endpoint URIs to retrieve token, refresh token, validate token, get a list of scopes associated with token and redirect user to obtain authorization_code. It specifies URIs to retrieve userinfo and to perform OpenID discovery to retrieve JSON web tokens and JSON web keys.

EXAMPLES

```
create oauth-provider f5Provider { authentication-uri https://local.f5.com/f5-oauth2/v1/authorize auto-jwt-
config-name auto_jwt_f5Provider last-discovery-time 2017-06-09:07:48:57 openid-cfg-uri
https://f5-oauth.local/f5-oauth2/v1/.well-known/openid-configuration description "OAuth provider defines F5
```

AS" token-uri https://local.f5.com/f5-oauth2/v1/token token-validation-scope-uri https://local.f5.com/f5-oauth2/v1/validate trusted-ca-bundle ca-bundle.crt type f5 }
Creates the OAuth Provider named f5Provider of type f5, defines all endpoint URLs to the local Authorization Server.

delete oauth-provider f5Provider
Deletes the OAuth Provider named f5Provider from the system.

OPTIONS

allow-self-signed-jwk-cert
Specifies whether creating JWK config containing self-signed certificate is allowed. The default value is yes.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

authentication-uri
The endpoint URI that is used to redirect user for authentication in order to get authorization_code. This endpoint is used by OAuth Client agent, when grant type is configured to Authorization Code.

auto-jwt-config-name
Specifies the name of the auto-discovered JWT config.

description
Specifies a description for the component. The default is none.

ignore-expired-cert
Specifies whether the expired AS certificate enforcement is to be ignored. The default value is false.

last-discovery-time
Specifies the last time JWT config and JWK config were auto-discovered.

manual-jwt-config-name
Specifies the name of the manually created JWT config.

max-json-nesting-layers
Specifies the maximum nesting layers of JSON responses that will be processed and stored. The default value is 8.

max-response-size
Specifies the size limit for all responses from all endpoints of the provider. The default value is 128kb.

openid-cfg-uri
The endpoint URI that is used to pull JSON web tokens and JSON web keys from Authorization server. This endpoint is used by OAuth Client agent.

save-json-payload
Specifies whether the entire JSON text is to be saved in a session variable. The default value is false.

partition
Displays the partition within which the component resides. The default is Common.

token-uri
The endpoint URI that is used to retrieve an access_token. This endpoint is used by OAuth Client agent.

token-validation-scope-uri
The endpoint URI that is used by OAuth Scope agent, in order to retrieve a list of scopes associated with an access_token. The same URI is used to validate an access_token by OAuth Client agent.

trusted-ca-bundle
Specifies the trusted CA bundle for AS certificate that is used for autodiscovery of JSON web tokens and JSON web keys.

type The type of the provider.

use-auto-jwt-config
Specifies whether the OAuth Provider uses auto-discovered JWT config specified in auto-jwt-config-name or manually created JWT config specified in manual-jwt-config-name. The default value is true.

userinfo-request-uri
The endpoint URI that is used to request userinfo information. This endpoint is used by OAuth Scope agent.

custom The provider is using a custom Authorization Server.
azure-ad The provider is using an AzureAD Authorization Server.
azure-ad-b2c The provider is using an AzureAD-B2C Authorization Server.
f5 The provider is using an F5 Authorization Server.
facebook The provider is Facebook Authorization Server.
google The provider is Google Authorization Server.
okta The provider is using an Okta Authorization Server.
ping The provider is Ping Identity Authorization Server.
Default value for provider type is f5.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016, 2017. All rights reserved.

BIG-IP 2018-01-18 apm aaa oauth-provider(1)

apm aaa oauth-request

NAME

oauth-request - Manages an OAuth Request.

MODULE

apm aaa

SYNTAX

Configure the oauth-request component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

```
create oauth-request [name]
modify oauth-request [name]
options:
  app-service [[string] | none]
  description [[string] | none]
  headers [add | delete | modify | replace-all-with] {
    [name] {
      value [value]
    }
  }
  method [get | post]
  parameters [add | delete | modify | replace-all-with] {
    [name] {
      type [parameter type]
      value [[string] | none]
    }
  }
  type [request type]
  uri [[string] | none]
```

```
edit oauth-request [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list oauth-request
list oauth-request [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete oauth-request [name]
```

DESCRIPTION

You can use the oauth-request component to manage an OAuth Request. The OAuth Request is an HTTP request that is used during communication between the BIG-IP system and an OAuth Authorization Server (AS). Different types of OAuth Requests can be configured for both OAuth Client and OAuth Scope agents.

EXAMPLES

```
create oauth-request F5AuthRedirectRequest { description "F5 Authentication Redirect request" method get
parameters add { client_id { type client-id } redirect_uri { type redirect-uri } response_type { value "code"
} } type auth-redirect-request }
Creates the OAuth Request named F5AuthRedirectRequest of type auth-redirect-request, sets HTTP method to get and specifies the list of GET parameters to be sent: client-id, redirect_uri, response_type.
```

```
delete oauth-request F5AuthRedirectRequest
Deletes the OAuth Request named F5AuthRedirectRequest from the system.
```

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

Specifies a description for the component. The default is none.

headers

Adds, deletes, or replaces a set of headers, by specifying a header name and value for each entry.

value

The value of the header.

method

Specifies the HTTP method for the OAuth Request. The options are:

get Configures the system to make HTTP request using GET method.

post Configures the system to make HTTP request using POST method.

parameters

Adds, deletes, or replaces a set of parameters, by specifying a parameter name for each entry. You can configure the following options for each parameter:

type The type of the parameter. For a custom type of parameter, you must provide a value. For other parameter types, the value is taken from other configurations. The options for the type of a parameter are:

access-token The value for the parameter is access_token. Value assigned from session variable session.oauth.client.access_token

client-id The value for this parameter type is the Client Id that is configured in the OAuth Server object.

client-secret The value for this parameter type is the Client Secret that is configured in the OAuth Server object.

grant-type The value for this parameter type is the Grant Type that is configured in the OAuth Client agent.

redirect-uri The value for this parameter type is the Redirect URI that is configured in the OAuth Client agent.

resource-server-id The value for this parameter is the Resource Server Id that is configured in the OAuth Server object.

resource-server-secret The value for this parameter is the Resource Server Secret that is configured in the OAuth Server object.

scope The value for this parameter is the Scope that is configured in the OAuth Client agent.

custom Custom parameter value; you can specify any custom value for the parameter.

Default value for parameter type is custom.

value

The value of the parameter. A value is required for parameters of type custom only.

partition

Displays the partition within which the component resides. The default is Common.

type Type of the request. The options for the type of a request are:

auth-redirect-request The Authentication Redirect request. This type of request is used to redirect user to an Authorization Server, when OAuth Client agent is configured to use "Authorization Code" grant type.

token-request The Token request. This type of request is used to access an Authorization Server in order to obtain an access_token or exchange an authorization_code for an access_token.

token-refresh-request The Refresh Token request. This type of request is used to refresh an expired access_token.

token-revocation-request The Revocation request. This type of request is used to revoke an access_token.

validation-scopes-request The Validation and Scopes request. This type of request is used in OAuth Client agent to validate an existing token. The same type of request is used to get a list of scopes associated with an existing token.

scope-data-request The Scope Data request. This type of request is used to obtain additional information from an Authorization Server.

Default value for request type is scope-data-request.

uri Request URI. This option is required for request of type scope-data-request only. All other types of requests use endpoint URIs configured at OAuth Provider component.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2017-01-20 apm aaa oauth-request(1)

apm aaa oauth-server

NAME

oauth-server - Manages an OAuth Server.

MODULE

apm aaa

SYNTAX

Configure the oauth-server component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create oauth-server [name]

modify oauth-server [name]

options:

app-service [[string] | none]

client-id [string]

client-secret [[string] | none]

client-serverssl-profile-name [name]

dns-resolver-name [name]

mode [client | rs | client-rs]

provider-name [name]

resource-server-id [string]

resource-server-secret [[string] | none]

resource-serverssl-profile-name [name]

rules [[string] | none]

token-validation-interval [[integer] | none]

edit oauth-server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list oauth-server

list oauth-server [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete oauth-server [name]

DESCRIPTION

You can use the oauth-server component to manage an OAuth Server. The OAuth Server specifies the configuration of an OAuth Authorization server for use by the OAuth Client or OAuth Scope agents.

EXAMPLES

```
create oauth-server f5Server { provider-name Google mode client client-id myClientId client-secret
e939e21ead60c0406341c9be587a005056890213d480f456 client-serverssl-profile-name serverssl dns-resolver-name
myResolver}
```

Creates the OAuth Server named f5Server and defines all required options. In this example, the BIG-IP system is supposed to only acquire an access_token from Google. The server mode is set to client and resource server credentials are not needed.

```
delete oauth-server f5Server
```

Deletes the OAuth Server named f5Server from the system.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

client-id

Specifies the client application ID. The client application must be configured before configuring the OAuth Server on the BIG-IP system.

client-secret

Specifies the client application secret. The client application must be configured at the authorization server before configuring the OAuth Server on the BIG-IP system.

client-serverssl-profile-name

SSL profile to be used by the BIG-IP system when connecting to authorization server.

dns-resolver-name

DNS resolver object to be used by OAuth Server to resolve DNS names for endpoint URIs.

mode The mode of operation for the OAuth Server. The options for the mode of operation are:

client The OAuth Server can be used by OAuth Client agent only. In this mode, you do not need to specify Resource Server credentials.

rs The OAuth Server can be used by OAuth Scope agent only. In this mode, you do not need to specify Client Application credentials.

client-rs The OAuth Server can be used by either OAuth Client or OAuth Scope agent. Client Application credentials and Resource Server credentials are required.

partition

Displays the partition within which the component resides. The default is Common.

resource-server-id

Specifies the Resource Server ID. The Resource Server must be configured before configuring OAuth Server on the BIG-IP system.

resource-server-secret

Specifies the Resource Server Secret. The Resource Server must be configured before configuring OAuth Server on the BIG-IP system.

resource-serverssl-profile-name

SSL profile to be used by the BIG-IP system when connecting to resource server.

rules

The list of iRule events. You can apply an iRule event to modify a request or a response (except an authorization code request from the BIG-IP OAuth client to the OAuth authentication server).

token-validation-interval

Specifies the number of minutes that the token can remain valid. The token becomes invalid when this interval elapses or at the token expiry that the authentication server specifies, whichever is shorter. When the token expires, the subsession times out. (This setting applies only to a per-request policy).

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2017-01-20 apm aaa oauth-server(1)

apm aaa ocsdp

NAME

ocsdp - Configure Online Certificate System Protocol (OCSP) responder objects.

MODULE

apm aaa

SYNTAX

Configure the ocsdp component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create ocsdp [name]

modify ocsdp [name]

options:

allow-certs [true | false]

app-service [[string] | none]

ca-file (| none)

ca-path (| none)

cert-id-digest (sha1 | md5)

chain [true | false]

check-certs [true | false]

explicit-ocsdp [true | false]

ignore-aia [true | false]

intern [true | false]

location-specific [true | false]

nonce [true | false]

sign-digest (sha1 | md5)

sign-key (| none)

sign-key-passphrase (| none)

sign-other (| none)

signer (| none)

status-age

trust-other [true | false]

url (| none)

va-file (| none)

validity-period

verify [true | false]

verify-cert [true | false]

verify-other (| none)

verify-sig [true | false]

edit ocsdp | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list oosp

list oosp [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete oosp [name]

DESCRIPTION

To implement the SSL OOSP authentication module, create an OOSP responder object and assign it to the OOSP auth agent in your access policy.

OPTIONS

allow-certs

Specifies whether the addition of certificates to an OOSP request is enabled. The default is true.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

ca-file

Specifies the name of the certificate file object containing trusted CA certificates used to verify the signature on the OOSP response. The default is none.

ca-path

Specifies the path to the trusted CA certificates used to verify the signature on the OOSP response. The default is none.

cert-id-digest

The cert ID digest is part of the OOSP protocol. The OOSP client (in this case, the BIG-IP system) calculates the cert ID using a hash of the Issuer and serial number for the certificate that it is trying to verify. The options are:

sha1 Newer algorithm that provides a higher security level with a 160 bit hash length. This is the default.

md5 Older algorithm with a 128 bit hash length.

chain

Specifies whether the system constructs a chain from certificates in the OOSP response. The default is true.

check-certs

Specifies whether the LTM system makes additional checks to see if the signer's certificate is authorized to provide the necessary status information. Use this option only for testing purposes. The default is true.

explicit-oosp

Specifies whether the BIG-IP system explicitly trusts that the OOSP response signer's certificate is authorized for OOSP response signing. If the signer's certificate does not contain the OOSP signing extension, setting this option to true causes a response to be untrusted. The default is true.

ignore-aia

Specifies whether to ignore the URL contained in the certificate's AIA fields, and to always use the URL specified by the responder instead. The default is false.

intern

Specifies whether to ignore certificates contained in an OOSP response when searching for the signer's certificate. When you set this option to true, you must also specify the signer's certificate using either the verify-other or va-file option. The default is true.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies a unique name for the component. This option is required.

nonce

Specifies whether a nonce will be sent in an OOSP request. When set to false, the request is sent without a nonce. The default is true.

partition

Displays the partition within which the OOSP responder object resides.

sign-digest

Specifies the algorithm (md5 or sha1) used to sign a request using a signing certificate and key. The default is sha1. If you use this option, you must also set the sign-key and sign-key-passphrase options.

sign-key

Specifies the key used to sign an OCSF request. If you use this option, you must also set the sign-digest and sign-key-passphrase options. The default is none.

sign-key-passphrase

Specifies the passphrase for the signing key. If you use this option, you must also set the sign-digest and sign-key options. The default is none.

sign-other

Specifies additional certificates to add to an OCSF request. The options are default.crt and ca-bundle.crt. The default is none.

signer

Specifies the certificate used to sign an OCSF request. If the certificate is specified but the key is not specified, then the private key is read from the same file as the certificate. If neither the certificate nor the key is specified, then the request is not signed. If the certificate is not specified and the key is specified, then the configuration is considered to be invalid. The default is none.

status-age

Specifies the amount of time (in seconds) to compare to the notBefore value of a status response. Use this option only when a status response does not include the notAfter field. The default is 0 (zero).

trust-other

Specifies whether the BIG-IP system trusts the certificates specified using the verify-other option. The default is false.

url Specifies the URL used to contact the OCSF service on the responder. This option is required. The default is none.

va-file

Specifies the name of the file containing explicitly-trusted responder certificates. Use this option when the responder is not covered by the certificates already loaded into the responder's CA store. The default is none.

validity-period

Specifies an acceptable error range in seconds. Use this option when the OCSF responder clock and a client clock are not synchronized, which could cause a certificate status check to fail. This value must be a positive number. This option is required. The default is 300.

verify

Specifies whether verification of an OCSF response signature or the nonce values is enabled. Use this option only for debugging purposes. The default is true.

verify-cert

Specifies whether the BIG-IP system verifies the certificate in the OCSF response. The default is true.

verify-other

Specifies the name of the file used to search for an OCSF response signing certificate when the certificate has been omitted from the response. The default is none.

verify-sig

Specifies whether the BIG-IP system checks the signature on the OCSF response. Use this option only for testing purposes. The default is true.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2016-01-07 apm aaa ocsp(1)

apm aaa okta-connector

NAME

okta-connector - stores the Okta API parameters.

MODULE

apm aaa

SYNTAX

Configure the okta-connector object within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create/modify okta-connector [name]

options:

domain [string]

token [string]
transport [name]

edit okta-connector [[[name] | [glob] | [regex]] ...]

options:
all-properties
non-default-properties

DISPLAY

list okta-connector

list okta-connector [[[name] | [glob] | [regex]] ...]

options:
all-properties
app-service
non-default-properties
one-line
partition

DELETE

delete okta-connector [name]

DESCRIPTION

okta-connector is used by Okta MFA agent to communicate with Okta service.

EXAMPLES

create okta-connector OktaConnectMFA domain dev-678901.okta.com token api-token transport
YourHttpConnectorTransport
Creates the Okta Connector object named OktaConnectMFA.

delete okta-connector OktaConnectMFA
Deletes the Okta Connector object named OktaConnectMFA from the system.

OPTIONS

name Specifies the name for the Okta Connector object. This setting is required.

domain
Specifies the Okta organization domain. This setting is required.

token
Specifies the Okta API token. This setting is required.

transport
Specifies the HTTP Connector Transport object. This setting is required.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2020. All rights reserved.

BIG-IP 2020-02-25 apm aaa okta-connector(1)

apm aaa radius

NAME

radius - Manages an AAA RADIUS server.

MODULE

apm aaa

SYNTAX

Configure the radius component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create radius [name]

modify radius [name]

options:

acct-port [integer]
address [[ip addr] | none]
auth-port [integer]
app-service [[string] | none]
description [[string] | none]
mode [acct | auth | both]
nas-ip-address [[ip addr] | none]
nas-ipv6-address [[ip addr] | none]
location-specific [true | false]
pool [[string] | none]

retries [integer]
secret [string]
service-type [default | login | framed | callback-login | callback-framed | outbound | administrative | nas-prompt | authenticate-only | callba
timeout [integer]
use-pool [enabled | disabled]

edit radius [[glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list radius

list radius [[[name] | [glob] | [regex]] ...]

options:

all-properties
app-service
non-default-properties
one-line
partition

DELETE

delete radius [name]

DESCRIPTION

You can use the radius component to create and manage an AAA RADIUS server.

EXAMPLES

```
create rad_auth { address 172.30.6.144 secret "test" use-pool "disabled" }
```

Creates the AAA RADIUS server named rad_auth that has an IP address of 172.30.6.144 and has a shared secret of test.

```
delete radius MyRadiusServer
```

Deletes the AAA RADIUS server named MyRadiusServer from the system.

OPTIONS

acct-port

Specifies the port number of the external AAA RADIUS accounting server. The default is radius-acct.

address

Specifies the IP address of the AAA RADIUS server. This option is required.

auth-port

Specifies the port number for the service. The default is radius. This option is required.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application can modify or delete the object.

description

Specifies a unique description for the AAA RADIUS server. The default is none.

mode Specifies the configuration mode you want to use for RADIUS authentication. Note that you cannot modify the mode once you create the server. The options are:

acct Configures the system to perform only RADIUS accounting. Use this option to pass accounting information about your users to the external RADIUS accounting server.

auth Configures the system to perform only RADIUS authentication. Use this option to authenticate your users through a RADIUS server.

both Configures the system to perform both RADIUS authentication and RADIUS accounting simultaneously.

[name]

Specifies the name of an AAA RADIUS server. This option is required.

nas-ip-address

Specifies an IP address as RADIUS attribute 4 that you can configure without changing the source IP address in the IP header of the RADIUS packets. Use this option in situations where you are using an NAS cluster to be recognized as a single RADIUS client.

nas-ipv6-address

Specifies an IPv6 address as RADIUS attribute 4 that you can configure without changing the source IP address in the IP header of the RADIUS packets. Use this option in situations where you are using an NAS cluster to be recognized as a single RADIUS client.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

partition

Displays the partition within which the component resides.

pool Specifies the name of the pool to which this server belongs. The default is none.

retries

Specifies the number of times the BIG-IP system tries to make a connection to the RADIUS AAA server after the first attempt fails. The default is 3.

secret

Specifies the shared secret password of the AAA RADIUS server. This option is required.

service-type

Specifies the type of service used for the RADIUS server. The default is default, which behaves as authenticate-only.

timeout

Specifies a timeout interval (in seconds) for the AAA RADIUS server after which the server closes a connection. The default is 5.

use-pool

Enables or disables the use of the pool specified using the pool option. The default is none.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm aaa radius(1)

apm aaa saml-idp-automation

NAME

saml-idp-automation - Specify SAML IdP automation configuration used to automate creation and management of 'IdP Connectors' from the remotely published metadata file(s).

MODULE

apm aaa

SYNTAX

Configure the saml-idp-automation component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

```
create saml-idp-automation [name]
modify saml-idp-automation [name]
options:
  aaa-saml-server [string]
  app-service [[string] | none]
  connection-properties [add | delete | modify | none | replace-all-with] {
    name [string] {
      app-service [[string] | none]
      dns-resolver-name [[string] | none]
      serverssl-profile-name [[string] | none]
    }
  }
  description [[string] | none]
  frequency [integer]
  idp-matching-source [string]
  idp-obj-name-tag [string]
  metadata-matching-tag [string]
  metadata-urls {
    [string]
  }
}
```

```
edit saml-idp-automation [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list saml-idp-automation
list saml-idp-automation [ [ [name] | [glob] | [regex] ] ... ]
show running-config saml-idp-automation
show running-config saml-idp-automation [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
```

partition

DELETE

delete saml-idp-automation [name]

DESCRIPTION

You can use saml-idp-automation to create and manage SAML IdP automation objects that are used to automate creation and management of 'IdP Connectors' from the remotely published metadata files.

EXAMPLES

```
create saml-idp-automation my_idp_automation1 { aaa-saml-server my_saml_sp frequency 60 idp-matching-source
"%{session.server.idpname}" metadata-matching-tag IdpName idp-obj-name-tag displayname metadata-urls add {
https://f5.com/metadata.xml } connection-properties add { cp1 { dns-resolver-name myResolver serverssl-
profile-name serverssl } } }
```

Creates a SAML IdP automation object named my_idp_automation1 bound to a SAML SP service my_saml_sp with frequency set to 60 minutes, idp-matching-source as %{session.server.idpname}, metadata-matching-tag as IdpName, idp-obj-name-tag as displayname, one entry for metadata-url as https://f5.com/metadata.xml and connection-properties with dns-resolver-name as myResolver and serverssl-profile-name as serverssl.

```
list saml-idp-automation
```

Displays a list of SAML IdP automation objects.

```
delete saml-idp-automation my_idp_automation1
```

Deletes the my_idp_automation1 SAML IdP automation object.

OPTIONS

```
aaa-saml-server
```

Specifies the AAA SAML server to which the IdP connectors created by this automation are bound.

```
app-service
```

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

```
connection-properties
```

Specifies the connection properties for fetching the metadata files. dns-resolver-name specifies the DNS resolver object to be used and serverssl-profile-name specifies the SSL profile to be used by the BIG-IP system when connecting to the server. Both DNS resolver and SSL profile should be configured if metadata files are located behind an SSL protected endpoint.

```
description
```

Specifies the description for the IdP automation object.

```
frequency
```

The frequency in minutes at which APM polls the IdP metadata files and updates the IdP connectors and bindings to the specified AAA SAML server. The default value is 60.

```
idp-matching-source
```

Specifies the selection criteria for IdP connectors. It must be in session variable format. It is used in configuration as a 'matching source' when binding created IdP connectors to configured AAA SAML server. At runtime, the value of this session variable is compared to metadata-matching-tag to determine which IdP connector is used to authenticate user.

```
metadata-matching-tag
```

This value is used in combination with idp-matching-source. It is used in configuration as a 'matching value' when binding created IdP connectors to configured AAA SAML server. At runtime, this value is compared against the value of session variable idp-matching-source to determine which IdP connector is used to authenticate user.

```
idp-obj-name-tag
```

Specifies the name of a tag within the metadata file that contains a value that APM includes in the names of the created IdP connectors.

```
metadata-urls
```

Specifies a list of one or more URLs containing the metadata files.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016, 2017. All rights reserved.

BIG-IP 2017-07-27 apm aaa saml-idp-automation(1)

apm aaa saml-idp-connector

NAME

saml-idp-connector - Specify saml idp connector configuration used for SAML authentication.

MODULE

apm aaa

SYNTAX

Configure the saml-idp-connector component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create saml-idp-connector [name]

modify saml-idp-connector [name]

options:

app-service [[string] | none]
artifact-resolution-service-addr [IP address]
artifact-resolution-service-port [integer]
artifact-resolution-service-url [[string] | none]
basic-auth-password [[string] | none]
basic-auth-username [[string] | none]
description [[string] | none]
entity-id [string]
identity-location [attribute | subject]
identity-location-attribute [[string] | none]
idp-certificate [[string] | none]
import-metadata [[metadata-file] | none]
location-specific [true | false]
metadata-cert [[string] | none]
name-qualifier [[string] | none]
serverssl-profile-name [profile name | none]
sign-artifact-resolution-rq [true | false]
single-logout-binding
single-logout-response-uri [[string] | none]
single-logout-uri [[string] | none]
sso-binding [http-post | http-redirect]
sso-uri [[string] | none]
want-authn-request-signed [true | false]
want-detached-signature [true | false]

edit saml-idp-connector [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list saml-idp-connector

list saml-idp-connector [[[name] | [glob] | [regex]] ...]

show running-config saml-idp-connector

show running-config saml-idp-connector [[[name] | [glob] | [regex]] ...]

options:

all-properties
app-service
non-default-properties
one-line
partition

DELETE

delete saml-idp-connector [name]

DESCRIPTION

You can use the saml-idp-connector to create and manage saml idp connectors.

EXAMPLES

```
create saml-idp-connector my_idp_connector { import-metadata /shared/tmp/meta_data_idp.xml }
```

Creates saml idp connector named my_idp_connector from metadata. In this example "/shared/tmp/meta_data_idp.xml" is a file containing saml identity provider metadata.

```
create saml-idp-connector my_idp_connector1 { entity-id "https://www.secureauth.com/dom1" identity-location
```

```
subject sso-binding http-post sso-uri "https://www.secureauth.com/dom1/acs/" idp-certificate my_company.crt }
```

Creates a saml idp connector named my_idp_connector1 with certificate "my_company.crt" with identity-location "subject".

```
list saml-idp-connector
```

Displays a list of saml idp connectors.

```
delete saml-idp-connector my_idp_connector
```

Deletes the my_idp_connector saml idp connector.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

artifact-resolution-service-addr

Specifies the IP address that this BIG-IP as SP will use to connect to the IdP artifact resolution service. Value must be a valid IPv4 or IPv6 address.

`artifact-resolution-service-port`

Specifies the port that this BIG-IP as SP will use to connect to the IdP artifact resolution service.

`artifact-resolution-service-url`

Specifies the URI of the IdP artifact resolution service. The URI must include protocol, hostname, and full path.

`basic-auth-password`

Specifies the password for basic authentication. When configured, basic authentication is used for the artifact resolve request sent to the IdP.

`basic-auth-username`

Specifies username for basic authentication. When configured, basic authentication is used for the artifact resolve request sent to the IdP.

`description`

Specifies a unique description for the saml idp connector. The default is none.

`entity-id`

Specifies unique URI to represent the IdP pointed by idp connector.

`identity-location`

Specifies location of user identity inside SAML assertion. It can be either one of the attributes or the subject.

`identity-location-attribute`

If the location of user identity is set to attribute then attribute name should be specified as part of this attribute.

`idp-certificate`

This is IdP's certificate and is used by BIG-IP as SP to verify the signature of the assertion.

`import-metadata`

This attribute specifies the metadata file from an external IdP system used for creating idp connector object.

For example: `create saml-idp-connector my_idp_connector { import-metadata /shared/tmp/meta_data_idp.xml}`

`location-specific`

Objects of this class might have location specific attribute(s). Admin can indicate if object is location specific by setting it to true.

`metadata-cert`

This specifies the certificate to use to verify the signature of metadata imported from a file.

For example: `create saml-idp-connector my_idp_connector2 {import-metadata /shared/tmp/meta_data_signed_idp.xml metadata-cert default.crt}`

`name-qualifier`

Specifies the security or administrative domain of the external IdP. This value usually matches IdP Entity ID.

`serverssl-profile-name`

Specifies the SSL profile used when this BIG-IP as SP connects to the IdP artifact resolution service.

`sign-artifact-resolution-rq`

Specifies whether the IdP requires artifact resolve requests to be signed. Default value is true.

`single-logout-binding`

This attribute is reserved for future functionality.

`single-logout-response-uri`

A URI where this BIG-IP as SP will send single logout (SLO) responses.

`single-logout-uri`

A URI where this BIG-IP as SP will send single logout (SLO) requests.

`sso-binding`

This specifies the method the IdP uses to receive authentication request from BIG-IP as SP. Default value is http-post

`sso-uri`

This specifies the URL of IdP's SSO service where BIG-IP as SP sends an authentication request to IdP.

`want-authn-request-signed`

This property specifies whether IdP requires signed authentication request. Set it to true if this BIG-IP as SP is required to send signed authentication request to IdP. Default value is false.

`want-detached-signature`

This property specifies signature type for messages sent by BIG-IP via HTTP Redirect binding. To use detached signatures set this property to true. Enveloped signatures are used by default.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

apm aaa saml

NAME

saml - Specify a SAML server configuration used for authentication.

MODULE

apm aaa

SYNTAX

Configure the saml component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

```
create saml [name]
modify saml [name]
options:
  app-service [[string] | none]
  assertion-consumer-binding [http-artifact | http-post]
  attribute-consuming-services [add | delete | modify | none | replace-all-with] {
    [name] {
  attribute-consuming-service-index [integer]
    }
  }
  auth-context-class-list [[string] | none]
  auth-context-comparison-method [ better | exact | maximum | minimum ]
  auth-context-methods {
[string]
  }
  default-attribute-consuming-service [[string] | none]
  description [[string] | none]
  entity-id [string]
  force-authn [true | false]
  export-metadata [ no-signing | with-signing ]
  idp-connectors [add | delete | modify | none | replace-all-with] {
    [name] {
  idp-matching-source [[string] | none]
  idp-matching-value [[string] | none]
    }
  }
  is-authn-request-signed [true | false]
  location-specific [true | false]
  metadata-cert [[string] | none]
  metadata-file [[string] | none]
  metadata-signkey [[string] | none]
  name-id-policy-allow-create [true | false]
  name-id-policy-format [[string] | none]
  name-id-policy-sp-name-qualifier [[string] | none]
  provider-name [[string] | none]
  relay-state [[string] | none]
  sp-certificate [[string] | none]
  sp-host [[string] | none]
  sp-scheme [http | https]
  sp-signkey [[string] | none]
  want-assertion-encrypted [true | false]
  want-assertion-signed [true | false]
```

```
edit saml [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
  all-properties
  non-default-properties
```

DISPLAY

```
list saml
list saml [ [ [name] | [glob] | [regex] ] ... ]
show running-config saml
show running-config saml [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete saml [name]
```

DESCRIPTION

You can use the saml component to create and manage saml aaa servers.

EXAMPLES

```
create saml my_saml_server { entity-id "https://spvs1.mycompany.com/id" want-assertion-signed true want-assertion-encrypted false is-authn-request-signed true sp-certificate my_company.crt sp-signkey my_company.key }
```

Creates a SAML authentication server named `my_saml_server` with certificate `my_company.crt` and key `my_company.key` and security options requiring signed assertion and want to send signed authentication request.

```
list saml
```

Displays a list of aaa saml servers.

```
delete saml my_saml_server
```

Deletes the `my_saml_server` aaa saml server.

OPTIONS

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`assertion-consumer-binding`

Specifies method this BIG-IP as SP uses to receive assertions. Default value is `http-post`.

`attribute-consuming-services`

Add one or more attribute consuming services to this SP service. Each attribute consuming service is mapped to a unique `attribute-consuming-service-index`. The attribute consuming services added for this SP will be part of the metadata for the SP that can be exported and shared with IdP.

For example:

The following command associates two attribute consuming services to an SP and maps the first service to index 1 and the second service to index 2.

```
modify saml my_saml_server attribute-consuming-services add { my_atcs1 { attribute-consuming-service-index 1 } my_atcs2 { attribute-
```

`auth-context-class-list`

Specifies an ordered list of authentication context classes. The BIG-IP as SP uses this list to validate the authentication context (in the assertion from the IdP) against locally configured context methods (`auth-context-methods`) using the specified comparison method (`auth-context-comparison-method`).

This property is required if you use a comparison method (`auth-context-comparison-method`) other than the default ('exact'). You can specify any `auth-context-class-list` list that you have configured on the BIG-IP system. Or, you can specify the predefined `auth-context-class-list` list (`authentication_contexts_list`) that the BIG-IP system provides.

`auth-context-comparison-method`

Specifies the comparison method that the IdP must use to evaluate the requested context classes `auth-context-methods`, one of "exact", "minimum", "maximum", or "better". The default is exact. If non-default comparison method is configured, all context classes from `auth-context-methods` must be present in the configured priority list of classes `auth-context-class-list`.

`auth-context-methods`

Specifies a list of authentication context classes that this BIG-IP as SP will request from an IdP. As a response, the IdP must return an assertion containing one of the requested authentication contexts. Each value can be a session variable if the comparison method is set to 'exact', which is the default value.

`default-attribute-consuming-service`

Specifies one of the attribute consuming services associated with this SP as default service. The metadata for the SP will flag specified service as default.

`description`

Specifies a unique description for the server. The default is none.

`entity-id`

Specifies a unique identifier for BIG-IP as SP. Typically 'entity-id' is a URI that points to the BIG-IP virtual server that is going to act as SAML SP. In case 'entity-id' is not a valid URL, the `sp-host` attribute is required. Examples of valid configuration include "https://mycompany-sp", "sp:my:company", and "sp.my.company.com".

`force-authn`

If enabled, this BIG-IP as SP requests the IdP to authenticate the principal directly rather than rely on a previous security context.

`export-metadata`

You can simplify SAML configuration using metadata files. When you use BIG-IP as an SP, you can export metadata for an SP to a file. Then you can use the file to configure SP metadata on an IdP system by importing the file or using the information in the file to configure the SP. You can choose to sign metadata while exporting it for better security.

For example:

1. Exporting metadata with signing. This requires `metadata-cert` and `metadata-signkey` files.

```
modify saml aaa_obj {export-metadata with-signing metadata-file /shared/sp_signed_metadata.xml metadata-cert default.crt metadata-si
```

2. Exporting metadata with no signing.

```
modify saml aaa_obj {export-metadata no-signing metadata-file /shared/sp_metadata.xml}
```

idp-connectors

Add one or more IdP connectors to this SP service. BIG-IP SP redirect users to associated IdPs for authentication. If more IdP connectors associated with the SP, BIG-IP SP selects one of the IdP based on the specified selection criteria.

For example:

1. The following command associates one IdP connect to an SP

```
modify saml my_saml_server idp-connectors add { my_idp_connector1 }
```

2. Following associates multiple IdP connectors to SP with selection criteria based on landing URI. If the landing URI is /google, the user is sent to IdP as specified by my_idp_connector_google_app and if the landing URI is /salesforce, the user is sent to IdP as specified by my_idp_connector_for_salesforce.

```
modify saml my_saml_server idp-connectors add { my_idp_connector_google_app { idp-matching-source "%{session.server.landinguri}" i
```

is-authn-request-signed

This property specifies whether the SP signs authentication requests while sending them to the IdP. Set it to true if this BIG-IP SP should sign authentication requests. The default value for this is false.

location-specific

Objects of this class might have location specific attributes. Admin can indicate if object is location specific by setting it to true.

metadata-cert

Specifies the certificate with public key of the key pair used in signing the metadata. See export-metadata for more information on metadata export functionality. This is the certificate to be included in signed metadata when we export metadata. This might or might not be SP certificate.

metadata-file

Specifies the file to which metadata is saved. See export-metadata for more information on metadata export functionality.

metadata-signkey

Specifies the key that is used to sign SP's metadata. See export-metadata for more information on metadata export functionality.

name-id-policy-allow-create

A Boolean value used to indicate whether external IdP is allowed, when processing requests from this BIG-IP as SP, to create a new identifier to represent the principal. Default value is false

name-id-policy-format

A URI reference representing the classification of string-based identifier information. For example, if a Service Provider (SP) initiates SSO by sending an AuthnRequest to the IDP with format "urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress", then the IdP response should contain subject identity in email format. This attribute can be a session variable.

name-id-policy-sp-name-qualifier

Optionally specifies that the assertion subject's identifier be returned in the namespace of an SP other than the requester, or in the namespace of a SAML affiliation group of SPs. This attribute can be a session variable.

relay-state

Specifies the value where the BIG-IP as SP redirects users after they are successfully authenticated and have been allowed by access policy. When BIG-IP receives the relay state from the IdP in addition to assertion, then it uses the value received from IdP to redirect the user to after authentication. Otherwise, BIG-IP uses the value from this configuration.

provider-name

Optionally specifies the human-readable name of this SAML SP for use by the identity provider.

sp-certificate

BIG-IP includes this certificate in the SAML SP metadata that you export. After the SAML SP metadata is imported on the IdP, the IdP can use this certificate to verify signed authentication request and to encrypt assertion.

sp-host

Hostname of this BIG-IP as SP. This attribute is required when "entity-id" is not a valid URL.

sp-scheme

Scheme used by this BIG-IP as SP. This attribute is only used when sp-host is not empty. Default value is https.

sp-signkey

This specifies the private key used to sign authentication requests if "is-authn-request-signed property" is set to true or to decrypt assertions when "want-assertion-encrypted" is set to true.

want-assertion-encrypted

This property specifies whether SP requires encrypted assertions. Set it to true if this BIG-IP SP requires encrypted assertions from the SAML IdP. The default value for this is false.

want-assertion-signed

This property specifies whether SP requires signed assertions. Set it to true if this BIG-IP SP requires signed assertions from the SAML IdP. The default value for this is true.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016, 2017. All rights reserved.

BIG-IP 2017-04-25 apm aaa saml(1)

apm aaa securid

NAME

securid - Manages an RSA SecurID authentication server.

MODULE

apm aaa

SYNTAX

Configure the securid component within the aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create securid [name]

modify securid [name]

options:

app-service [[string] | none]

config-files [[string] | none]

description [[string] | none]

location-specific [true | false]

source-ip [ip addr]

edit securid [[glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list securid

list securid [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete securid [name]

DESCRIPTION

You can use the securid component to create and manage an RSA SecurID authentication server.

EXAMPLES

```
create securid mySecuridServer { config-files add { sdconf.rec { local-path /shared/tmp/1 } } source-ip 172.31.54.138 }
```

Creates the mySecuridServer AAA RSA SecurID server.

```
list securid all
```

Displays a list of AAA RSA SecurID servers on the system.

```
delete securid mySecuridServer
```

Deletes the mySecuridServer AAA RSA SecurID server from the system.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

config-files

Specifies which files to use for SecurID authentication. Upload a copy of the sdconf.rec file from your RSA Authentication Manager server.

description

Specifies a description for the configuration file you are uploading.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

source-ip

Specifies the source IP address of the RSA SecurID agent. This option is required when authenticating to the RSA Authentication Manager server.

partition

Displays the partition within which the component resides.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 apm aaa securid(1)

apm aaa tacacsplus

NAME

tacacsplus - Configure a TACACS+ server for implementing remote TACACS+-based client authentication.

MODULE

apm aaa

SYNTAX

Configure the tacacsplus component within the apm aaa module using the syntax shown in the following sections.

CREATE/MODIFY

create tacacsplus

modify tacacsplus

options:

address [ip addr]

auth-service [arap | enable | fwproxy | login | nasi | none | ppp | pt | rcmd | x25]

auth-type [arap | ascii | chap | mschap | pap]

app-service [[string] | none]

description [[string] | none]

encrypt [enabled | disabled]

location-specific [true | false]

pool [[string] | none]

port [[string] | none]

priv-lvl [max | min | user]

protocol [atalk | deccp | ftp | http | ip | ipx | lat | lcp | osicp | pad | rlogin | telnet | tn3270 | unknown | vines | vpdn | xremote]

secret [[string] | none]

service [none | arap | connection | firewall | ppp | shell | slip | system | tty-daemon]

use-pool [[string] | none]

edit tacacsplus [[glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tacacsplus

list tacacsplus [[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete tacacsplus [name]

DESCRIPTION

You can use the tacacsplus component to create and manage a TACACS+ authentication server.

EXAMPLES

create tacacsplus mytacacs auth-service enable encrypt enabled

Creates a TACACS server named mytacacs with encryption enabled.

OPTIONS

address

Specifies the IP address of the TACACS+ server. This option is required.

auth-service

Specifies the name of the service that the user is requesting to be authenticated to use. This enables the TACACS+ server to behave differently for different types of authentication requests. This option is required.

auth-type

Specifies the type of authentication to be used for authenticating the user.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

Specifies a unique description for the component. The default is none.

encrypt

Enables or disables encryption of TACACS+ packets. Recommended for normal use. The default is enabled.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies the name of an AAA TACACS+ server. This option is required.

partition

Displays the partition within which the component resides.

pool Specifies the name of the pool to which this server belongs. The default is none.

port Specifies the port number of the server. The default is 49.

priv-lvl

Specifies the privilege level at which the user is authenticating. The options are:

max

min This is the default.

user

protocol

Specifies the protocol associated with the value specified in the service option, which is a subset of the associated service being used for client authorization or system accounting. The default is unknown.

secret

Sets the secret key used to encrypt and decrypt packets sent or received from the server. This option is required.

service

use-pool

Enables or disables the use of the pool specified using the pool option. The default is none.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2014-10-27 apm aaa tacacsplus(1)

apm access-info

NAME

access-info - Shows session related information such as session id, client ip, logon user and access profile name.

MODULE

apm access-info

SYNTAX

Shows the access-info component using the syntax shown in the following sections.

DISPLAY

show access-info

options:

all-properties
save-to-file

show access-info all

show access-info logon-user username

show access-info client-ip ip-address

EXAMPLES

show access-info

Displays the session id, client ip and logon user.

show access-info logon-user username

Displays the session id, and the corresponding client ip.

show access-info client-ip ip-address

Displays the session id and the corresponding logon user.

show access-info all-properties

The output of this command looks like this:

```
session      session_id
partition    partition Name
status       status of the session such as established and active
logon-user   username
client-ip    client ip address
virtual-server virtual server name
profile-access-type profile access type
access-profile access profile name
start-time   start time of the session
expiration-time expiration time of the session
rx-bytes     session traffic bytes in
tx-bytes     session traffic bytes out
rx-packets   session traffic packets in
tx-packets   session traffic packets out
ingress-raw  ingress raw bytes
ingress-compressed ingress compressed bytes
egress-raw   egress raw bytes
egress-compressed egress compressed bytes
slot         slot number (applicable only in the cluster environment)
tmm          tmm number for the session
```

Displays the session id, client ip, logon user and other information of session as shown above.

OPTIONS

session

Specifies the session id.

partition

Specifies the virtual server partition name.

status

Specifies the access status and valid values are established or pending.

logon-user

Specifies the logged in user name for the session.

client-ip

Specifies the client IP from which the session is established.

virtual-server

Specifies the name of the virtual server to which the session is established.

profile-access-type

Specifies access profile type; examples of valid values are: all, ltm-apm, ssl-vpn, sso, swg-explicit and swg-transparent.

access-profile

Specifies the access profile name.

start-time

Specifies the start time of the session.

expiration-time

Specifies the expiration time of the session.

rx-bytes)

Specifies the traffic bytes in for the session.

tx-bytes)

Specifies the traffic bytes out for the session.

rx-packets

Specifies the traffic packets in for the session.

tx-packets

Specifies the traffic packets out for the session.

ingress-raw

Specifies the ingress raw bytes for the session.

ingress-compressed

Specifies the ingress compressed bytes for the session.

egress-raw

Specifies the egress raw bytes for the session.

egress-compressed

Specifies the egress compressed bytes for the session.

slot Specifies the slot number for the session. Applicable only in the cluster environment.

tmm Specifies the tmm number that holds the session.

SEE ALSO

apm session

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2016. All rights reserved.

BIG-IP 2017-07-28 apm access-info(1)

apm acl

NAME

acl - Manages an access control list (ACL).

MODULE

apm

SYNTAX

Configure the acl component within the apm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create acl [name]
```

```
modify acl [name]
```

options:

```
acl-order [integer]
```

```
app-service [[string] | none]
```

```
description [[string] | none]
```

```
entries {
```

```
{
```

options:

```
action [allow | continue | discard | reject | unspec]
```

```
dst-end-port [[service] | none]
```

```
dst-start-port [[service] | none]
```

```
dst-subnet [[ip addr] | [[ip addr] [mask]]]
```

```
host [[string] | none]
```

```
log [config | none | packet | summary | verbose]
```

```
paths [[string] | none]
```

```
protocol [integer]
```

```
scheme [any | http | https]
```

```
src-end-port [[service] | none]
```

```
src-start-port [[service] | none]
```

```
src-subnet [[ip addr] | [[ip addr] [mask]]]
```

```
}
```

```
}
```

```
location-specific [true | false]
```

```
path-match-case [false | true]
```

```
type [dynamic | static]
```

DISPLAY

```
list acl
```

```
list acl [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
partition
```

show acl
show acl [[[name] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE
delete acl [name]

DESCRIPTION

You can use the acl component to configure a set of restrictions associated with a resource or favorite that defines access for users and groups.

EXAMPLES

```
create acl MyACL { acl-order 3 entries src-start-port ip default inet dst-end-port ip default inet action allow }
```

Creates the static access control list named MyACL that is the third ACL in the list of ACLs in the visual policy editor, and adds an access control entry that allows traffic using the default source IP address and the default destination IP address.

```
list acl all-properties
```

Displays a list of ACLs that includes the attributes of each ACL.

```
delete acl MyACL
```

Deletes the MyACL access control list.

OPTIONS

acl-order

Specifies the order of the access control entries in this access control list. This option is required.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

Describes the access control list.

entries

Configures an entry for an access control list.

action

Specifies the action that an access control list takes when this access control list entry is encountered. This option is required. You can specify one of the following actions:

allow

Allows traffic.

continue

Skips checking against the remaining access control list entries in this access control list, and continues evaluation at the next access control list.

discard

Drops packets silently.

reject

Drops a packet and sends TCP RST on TCP flows or proper ICMP messages on UDP flows. Silently drops a packet on other protocols.

dst-end-port

Specifies the destination IP address and network mask of the access control list entry. The default is 0.

dst-start-port

Specifies the source port or range of ports of the access control list entry.

dst-subnet

Specifies the destination subnet.

host Specifies the host name of the access control list entry.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

log Specifies the log level that is logged when actions of this type occur. Your options are:

config

Logs the configuration of a matched entry.

none Logs nothing. This is the default value.

packet

Logs a matched packet.

summary

Logs the name and entry number of a matched access control list and access control list entry.

verbose

Logs everything.

paths

Specifies an L7 access control list of matching URL paths.

protocol

Specifies the protocol number (TCP=6, UDP=17) of the access control list entry. The default is 0.

src-end-port

Specifies the source IP address and network mask of the access control list entry.

src-start-port

Specifies the source port or range of ports of the access control list entry.

src-subnet

Specifies the source subnet.

[name]

Specifies the name of the access control list. This setting is required.

partition

Displays the partition within which the object resides. The default is Common.

path-match-case

Indicates whether the path is case sensitive. The default is true.

type Specifies the type of access control list. The default is static. The available types are static and

dynamic.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2016-08-12 apm acl(1)

apm apm-avr-config

NAME

apm-avr-config - Configures AVR overview/statistics settings for Secure Web Gateway Functionality

MODULE

apm

SYNTAX

Configure an apm-avr-config component within the apm module using the syntax shown in the following sections.

MODIFY

The AVR Configuration consists of the following information: boolean flag to turn off data collection by AVR; Boolean flag to turn off sampling by AVR; Both are on by default. No create or delete function is allowed.

modify apm-avr-config apm-avr-config {

avr-collect-data [true | false]

avr-sampling [true | false]

}

DISPLAY

list apm-avr-config

list apm-avr-config [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DESCRIPTION

Configures AVR Configuration for SWG Statistical Reporting

EXAMPLES

modify apm-avr-config apm-avr-config { avr-collect-data false }

Modify apm-avr-config by setting avr collect data to false.

OPTIONS

avr-collect-data

Specifies whether data should be collected or not for statistical reporting.

avr-sampling

Specifies whether sampling should be turned on or off.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2014-02-18 apm apm-avr-config(1)

apm client image

NAME

image - Manages APM client software images.

MODULE

apm client

SYNTAX

Install, display information about, or delete a software image using the syntax in the following sections.

INSTALL

list image [name]

DISPLAY

list image

list image [[[name [/slot_id]] | [glob] | [regex]] ...]

options:

checksum

config-source

file-size

last-modified

one-line

DELETE

delete image [[[name] ...] | [all]]

DESCRIPTION

You can use the image component to install images, view information about available images, or delete unwanted images.

INSTALLING A SOFTWARE IMAGE

Before you begin installing an image, you must download the image file into the /shared/apm/images directory. You can find new software images at <http://downloads.f5.com>. We recommend downloading both the .iso file and the .md5 file. Download the file (or files) to your local machine, then transfer it to the /shared/apm/images directory on the BIG-IP. Use the Manager (GUI) interface to make this transfer, or quit tmsh to the Unix command line and use scp or a similar Unix command.

If you downloaded the .md5 file, you can use the Unix md5sum command to check the MD5 hash of the .iso file, and you can compare it to the contents of the .md5 file. They should match. If they do not, retry the download and/or transfer of the .iso file.

Use the install command with this component to install the .iso file.

Note: You use the slot_id option only for chassis systems and only when displaying the values for the options of a specific image. You do not use the slot_id option when installing or deleting an image, because these commands operate on all blades or the entire system.

EXAMPLES

```
list image apmclients-13.0.0-1914.0.iso
```

Displays information about the specified image, apmclients-13.0.0-1914.0.iso.

```
list image */1
```

Displays information about all of the images located on the first slot.

```
install image apmclients-13.0.0-1914.0.iso
```

Attempts to install the specified image, apmclients-13.0.0-1914.0.iso.

delete image apmclients-13.0.0-1914.0.iso

Removes the specified image, apmclients-13.0.0-1914.0.iso from /shared/apm/images directory.

OPTIONS

checksum

Displays the checksum of the image. You can use this option to verify the integrity of the image.

config-source

Displays the source of the image: "user" or "base".

file-size

Displays the size of the image file in megabytes.

last-modified

Displays the date the file was last modified.

name Specifies the name of the image that you want to install or delete.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

delete, glob, install, list, regex, tmsh, show

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2017. All rights reserved.

BIG-IP 2017-05-17 apm client image(1)

apm configuration captcha

NAME

captcha - Manages CAPTCHA version 2 configuration

MODULE

apm configuration

SYNTAX

Configure the captcha component within the configuration module using the syntax shown in the following sections.

CREATE/MODIFY

create captcha [name]

modify captcha [name]

options:

captcha-data-size [data-size-compact | data-size-normal]

captcha-data-theme [data-theme-dark | data-theme-light]

captcha-data-type [data-type-audio | data-type-image]

captcha-theme [theme-red | theme-white | theme-blackglass | theme-clean | theme-custom]

challenge-url [string]

description [string]

exposition-threshold [integer]

noscript-url [string]

private-key [hexadecimal string]

proceed-on-verification-error [false | true]

public-key [hexadecimal string]

secret [hexadecimal string]

site-key [hexadecimal string]

track-by-ip [false | true]

track-by-username [false | true]

verification-url [string]

edit captcha [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list captcha

list captcha [[[name] | [glob] | [regex]] ...]

options:

all-properties
app-service
non-default-properties
one-line
partition

DELETE
delete captcha [name]

DESCRIPTION
You can use the captcha component to create and manage CAPTCHA configuration.

EXAMPLES
create captcha mycaptcha {secret 123456789abcdef site-key fedcba987654321 }
Creates a CAPTCHA version 2 configuration named mycaptcha that uses site secret 123456789abcdef and site key fedcba987654321.

delete captcha mycaptcha
Deletes the CAPTCHA version 2 configuration named mycaptcha from the system.

OPTIONS
captcha-data-size
Specifies the size of the reCAPTCHA widget. The default is data-size-normal.

captcha-data-theme
Specifies the color theme of the reCAPTCHA widget. The default is data-theme-light.

captcha-data-type
Specifies the type of CAPTCHA to server. The default is data-type-image.

captcha-theme
This option is specific to reCAPTCHA v1 and is deprecated in version 13.0.0.

challenge-url
Specifies the URL of the service that provides the CAPTCHA challenge. The default is www.google.com/recaptcha/api.js.

description
Specifies a unique description for the CAPTCHA configuration.

exposition-threshold
Specifies the number of logon attempts to allow before issuing a CAPTCHA challenge. The default is 0.

noscript-url
Specifies the URL to use for obtaining the challenge if JavaScript is disabled. The default is www.google.com/recaptcha/api/fallback.

private-key
This option is specific to reCAPTCHA v1 and is deprecated in version 13.0.0.

proceed-on-verification-error
Specifies whether to allow user access when CAPTCHA verification cannot be completed for some reason. The default is true.

public-key
This option is specific to reCAPTCHA v1 and is deprecated in version 13.0.0.

secret
Specifies the secret provided by the CAPTCHA service provider. This option is required.

site-key
Specifies the site key provided by the CAPTCHA service provider. This option is required.

track-by-ip
Specifies whether to track logon failures using IP address. The default is true.

track-by-username
Specifies whether to track logon failures using username. The default is true.

verification-url
Specifies the URL of the service that verifies the response to the CAPTCHA challenge. The default is www.google.com/recaptcha/api/siteverify.

SEE ALSO
COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2016. All rights reserved.

BIG-IP 2016-11-01 apm configuration captcha(1)

apm epsec epsec-package

NAME

epsec-package - Manages an EPSEC package.

MODULE

apm epsec

SYNTAX

Configure the epsec-package component within the apm epsec module using the syntax shown in the following sections.

CREATE

create epsec-package [name]

options:

local-path [string]

server [[string] | none]

DISPLAY

list epsec-package

options:

all-properties

non-default-properties

recursive

list epsec-package [name]

INSTALL

install epsec-package [name]

options:

device-group [string]

DELETE

delete epsec-package [name]

DESCRIPTION

You can use the epsec-package component to create, install and manage an EPSEC package.

EXAMPLES

create epsec-package epsec_package_name local-path file_path

Creates an EPSEC package named epsec_package_name under the /Common/EPSEC/Upload folder from EPSEC image file located at file_path. Note: epsec_package_name must begin with the 'epsec' prefix.

list epsec-package

Displays a list of EPSEC packages under the specific folder. To list all EPSEC packages use the recursive option.

install epsec-package epsec_package_name

Installs the EPSEC package named epsec_package_name on this device. You cannot install a package from the /Common folder as it is a pre-installed package.

install epsec-package epsec_package_name device-group /Common/my_epsec_dg

Installs the EPSEC package named epsec_package_name on the devices in the device group /Common/my_epsec_dg. You cannot install a package from the /Common folder as it is a pre-installed package.

delete epsec-package epsec_package_name

Deletes the EPSEC package named epsec_package_name.

OPTIONS

[name]

Specifies the name of the EPSEC package. This option is required. Note: The name must begin with the 'epsec' prefix.

local-path

Specifies the local path of the EPSEC image file used with the CREATE command. This option is valid only with the CREATE command and is a required option.

device-group

Specifies the device group on which the package will be installed. This option is valid only with INSTALL command

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

apm epsec software-status

NAME

software-status - Displays the status of the EPSEC software installation.

MODULE

apm epsec

SYNTAX

Display information about the software-status component within the apm epsec module using the following syntax.

DISPLAY

show software-status

DESCRIPTION

You can use the software-status component to display the status of the EPSEC software installation, including the version of the EPSEC package being installed and the OESIS software version.

EXAMPLES

show software-status

Displays the status of the EPSEC software installation in a table.

OPTIONS

SEE ALSO

epsec-package

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-08-17 apm epsec software-status(1)

apm license

NAME

license - Shows the session information related to apm license.

MODULE

apm

SYNTAX

Displays the apm license information.

DISPLAY

show apm license

DESCRIPTION

APM module license is based on the session count depending on the platform. This module shows the total session information for access, ccu (connectivity), swg, swg limited and the currently used sessions for access, ccu, swg, swg limited. In addition, the statistics include the threshold percent for all the sessions.

EXAMPLES

show apm license

Displays the apm license usage information.

total access sessions

Total access sessions for BIG-IP. This number is based on the license and platform type.

current active sessions

The number of access sessions that are currently in use.

current established sessions

Total number of currently established sessions.

access sessions threshold percent

Access sessions threshold warning percent set by the user. The default is 75%.

total connectivity sessions

Total connectivity sessions (ccu) for BIG-IP. This number is based on the license and platform type.

current connectivity sessions

The number of connectivity sessions (ccu) that are currently in use.

connectivity sessions threshold percent

Connectivity sessions threshold warning percent set by the user. The default is 75%.

total swg sessions

Total swg sessions for BIG-IP. This number is based on the subscription-based swg license.

current swg sessions

The number of swg sessions that are currently in use.

swg sessions threshold percent

SWG sessions threshold warning percent set by the user. The default is 75%.

total swg limited sessions

Total swg limited sessions for BIG-IP. This number is based on the license and platform type.

current swg limited sessions

The number of swg limited sessions that are currently in use.

swg limited sessions threshold percent

SWG limited sessions threshold warning percent set by the user. The default is 75%.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014-2015. All rights reserved.

BIG-IP 2015-06-05 apm license(1)

apm log-setting

NAME

log-setting - Configures log configurations for various features in APM, such as URL Filter/Classification (URL Filter).

MODULE

apm

SYNTAX

Configure the log-setting component within the module using the syntax shown in the following sections.

CREATE/MODIFY

Consider log-setting as a container for log configurations belonging to different features. At this time URL Filter is the only feature with a log setting.

```
create log-setting [name]
modify log-setting [name]
options:
  description
  url-filters [add | delete | modify | replace-all-with] {
    [item name] {
  filter { log-allowed-url [true|false] log-blocked-url [true|false] }
  publisher [string]
    }
  }
  access [add | delete | modify | replace-all-with] {
    [item name] {
  log-level {
    access-control [emerg|alert|crit|err|warn|notice|info|debug]
    access-per-request [emerg|alert|crit|err|warn|notice|info|debug]
    apm-acl [emerg|alert|crit|err|warn|notice|info|debug]
    eca [emerg|alert|crit|err|warn|notice|info|debug]
    paa [emerg|alert|crit|err|warn|notice|info|debug]
    sso [emerg|alert|crit|err|warn|notice|info|debug]
    swg [emerg|alert|crit|err|warn|notice|info|debug]
  }
  publisher [string]
  }
}
```

DISPLAY

```
list log-setting
list log-setting [ [ [name] | [glob] | [regex] ] ... ]
options:
```

all-properties
non-default-properties
one-line
partition

DESCRIPTION

Configures a container for log configurations.

NOTE: Each container can enclose log configurations for many different features. However, each feature can only have one log configuration in a container.

NOTE: For the log configuration to take effect, the log-setting must be associated with an access profile (See man page for apm access profile).

NOTE: A log-setting container cannot be deleted if it is associated with an access profile.

EXAMPLES

```
create log-setting my-log-cfg
```

Creates a container without any log configuration.

```
create log-setting my-log-cfg url-filters add { my-urlf { filter { log-allowed-url true } publisher my-publisher } }
```

Creates a container with a log configuration for the URL Filter feature. At this version, URL Filter is the only feature with a log configuration.

```
modify log-setting my-log-cfg url-filters modify { my-urlf { publisher my-other-publisher } }
```

Modify the publisher of a log configuration.

```
modify log-setting my-log-cfg url-filters modify { my-urlf { filter { log-allowed-url false } } }
```

Modify the setting of a log filter

```
create log-setting my-log-cfg access add { my-access { publisher my-publisher } }
```

Creates a container with a log configuration for APM logging.

```
modify log-setting my-log-cfg access modify { my-access { log-level { access-control debug } } }
```

Modify the log level for module access-control.

OPTIONS

description

Specifies a unique description for the log-setting container.

url-filters

This is the list to store log configurations for the URL Filter feature.

item name

Specifies the name of the log configuration to be added to the list url-filters. Currently, the list supports only one item.

filter

Specifies the value for different log filters. In particular, URL Filter log configuration has two filters:

log-allowed-url [true|false]

log-blocked-url [true|false]

publisher

Specifies the publisher of the log configuration. See sys log-config publisher.

access

This is the list to store log configurations for APM logging.

log-level

This is the list of log level settings for different modules in the APM family.

SEE ALSO

apm profile access and sys log-config

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013-2014, 2016. All rights reserved.

BIG-IP 2016-04-21 apm log-setting(1)

apm ntlm machine-account

NAME

machine-account - Configures an APM NTLM machine account object.

MODULE

apm ntlm

SYNTAX

Configure the ntlm machine account using the syntax shown in the following sections.

CREATE/MODIFY

create machine-account [name]

options:

action [noop]

administrator-name [[string] | none]

administrator-password [[string] | none]

app-service [[string] | none]

domain-controller-fqdn [fqdn]

domain-fqdn [fqdn]

machine-account-name [[string] | none]

modify machine-account [name]

options:

action [change-password | noop]

app-service [[string] | none]

domain-controller-fqdn [fqdn]

edit machine-account [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list machine-account

list machine-account [[name] | [glob] | [regex]] ...]

DELETE

delete machine-account [name]

DESCRIPTION

You can use the machine-account component to configure a NTLM machine account.

EXAMPLES

```
create machine-account myaccount { machine-account-name "my_account_name" domain-fqdn "company.com" domain-controller-fqdn "server01.company.com" administrator-name "administrator" administrator-password "!My123Password" }
```

Creates a NTLM machine account named myaccount in the company.com domain, with domain controller server01.company.com, administrator name administrator and administrator password !My123Password.

```
list machine-account
```

Displays a list of all NTLM machine accounts created on the system.

```
delete machine-account myaccount
```

Deletes the NTLM machine account named myaccount the system.

OPTIONS

machine-account-name

Specifies the name of the machine account.

domain-fqdn

Specifies the Fully Qualified Domain Name. This setting is required.

domain-controller

Specifies the Fully Qualified Domain Name (FQDN) of the domain controller for the domain specified in the domain-fqdn option. The default is none.

administrator-name

Specifies the name of a user that has administrative permissions on an Active Directory server. This setting is required only when a new machine account is being created.

administrator-password

Specifies the password associated with administrator-name. This setting is required only when a new machine account is being created.

action

Specifies the action type. To change the machine account password, type this action: change-password else noop

change-password

Specifies that you want to change the machine account password.

noop Specifies "no operation performed". This is the default.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO
ntlm-auth

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 apm ntlm machine-account(1)

apm ntlm ntlm-auth

NAME
ntlm-auth - Configures an APM NTLM authentication object.

MODULE
apm ntlm

SYNTAX
Configure the ntlm-auth using the syntax shown in the following sections.

CREATE/MODIFY

```
create ntlm-auth [name]
options:
  app-service [[string] | none]
  dc-fqdn-list [add | delete | modify | replace-all-with] {
[[string]]
  }
  machine-account-name [[string] | none]
```

```
modify ntlm-auth [name]
options:
  app-service [[string] | none]
  dc-fqdn-list [add | delete | modify | replace-all-with] {
[[string]]
  }
  machine-account-name [[string] | none]
```

```
edit ntlm-auth [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list ntlm-auth
list ntlm-auth [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete ntlm-auth [name]
```

DESCRIPTION
You can use the ntlm-auth component to configure an NTLM authentication object.

EXAMPLES
create ntlm-auth myaccount { dc-fqdn-list add { server01.company.com } machine-account-name "my_account" }
Creates a NTLM authentication object named myaccount with machine account my_account, and the list of domain controllers specified by dc-fqdn-list

```
list ntlm-auth
Displays a list of all NTLM authentication objects created on the system.
```

```
delete ntlm-auth myaccount
Deletes the NTLM authentication object named myaccount from the system.
```

OPTIONS

- dc-fqdn-list
Specifies a list of Fully Qualified Domain Names (FQDNs) for the domain controllers to use for NTLM authentication.
- machine-account-name
Specifies the NTLM machine account object name to use for this NTLM authentication
- app-service
Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

machine-account

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012. All rights reserved.

BIG-IP 2014-05-27 apm ntlm ntlm-auth(1)

apm oauth db-instance

NAME

db-instance - Manages various OAuth database instances for this partition.

MODULE

apm oauth

SYNTAX

Configure the db-instance component within the oauth module using the following syntax.

CREATE/MODIFY

```
create db-instance [name]
```

```
modify db-instance [name]
```

options:

```
app-service [[string] | none]
```

```
description [[string] | none]
```

```
purge-frequency [daily | hourly | monthly | never | weekly]
```

```
purge-now
```

```
purge-time [string]
```

```
edit db-instance [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DISPLAY

```
list db-instance
```

```
list db-instance [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config db-instance
```

```
show running-config db-instance [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DELETE

```
delete db-instance [name]
```

DESCRIPTION

OAuth authorization server supports persistence for OAuth tokens and related data using on-disk databases. This data can be grouped into separate logical entities using database instances. The db-instance component can be used to manage such database instances within a partition.

Because disk size is a limited resource, expired/obsolete tokens need to be purged in order to make space for newly issued tokens. The db-instance component provides options to purge either periodically, using combination of purge-frequency and purge-time or on-demand, using purge-now. The default setting for purge-frequency is daily and for purge-time is 02:00 hours for a newly created database instance. When executed, database instance purging removes revoked, expired access tokens, refresh tokens, auth code and associated entries from the particular instance. Expired access tokens will not be removed if the reuse-access-token setting is enabled in the corresponding OAuth profile.

EXAMPLES

```
create db-instance myDbInstance {
  description "Sales Team"
  purge-frequency weekly
  purge-time "00:00"
}
```

Creates a database instance named myDbInstance that, when associated with an OAuth profile, stores tokens related to the OAuth profile in a separate database. This database instance is for the Sales team within the company. Stored data is purged weekly at 12AM.

```
list db-instance
```

Displays a list of all DB instances on the Authorization server.

```
delete db-instance myDbInstance
```

Deletes the database instance myDbInstance.

OPTIONS

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`db-name`

Specifies the actual name of the DB on BIG-IP. The value is unique for each DB instance and can be used to debug oauth DB interactions with BIG-IP storage.

`description`

Specifies a user-defined description for the database instance. The default value is none.

`[name]`

Specifies the name of the OAuth database instance. This setting is required.

`partition`

Displays the partition within which the component resides.

`purge-frequency`

Specifies the frequency at which data should be purged. The default value is daily. Other possible values are hourly, monthly, never and weekly.

`purge-now`

Indicates a request to purge the data right now.

`purge-time`

Specifies the time at which data should be purged. The default value is 02:00. HH:MM format must be used to specify time value.

SEE ALSO

`apm profile oauth`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2016. All rights reserved.

BIG-IP 2017-06-21 apm oauth db-instance(1)

apm oauth jwk-config

NAME

`jwk-config` - Manages JSON Web Keys to be used with Authorization Server/Client/Resource Server

MODULE

`apm oauth`

SYNTAX

Configure the `jwk-config` component within the `oauth` module using the following syntax.

CREATE/MODIFY

```
create jwk-config [name] modify jwk-config [name]
```

options:

`alg-type` [none | HS256 | HS384 | HS512 | RS256 | RS384 | RS512 | ES256 | ES384]

`app-service` [[string] | none]

`auto-generated` [enabled | disabled]

`cert` [certificate-name | none]

`cert-chain` [chain-name | none]

`cert-key` [key-name | none]

`cert-thumbprint-sha1` [[string] | none]

`cert-thumbprint-sha256` [[string] | none]

`curve` [[string] | none]

`include-x5c` [enabled | disabled]

`key-id` [[string] | none]

`key-type` [rsa | octet | elliptic-curve]

`key-use` [signing]

`modulus` [[string] | none]

`passphrase` [[string] | none]

`public-exponent` [[string] | none]

```
shared-secret [[string] | none]
use-client-secret [true | false]
x-coordinate [[string] | none]
y-coordinate [[string] | none]
```

```
edit jwk-config [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
one-line
```

DISPLAY

```
list jwk-config
```

```
list jwk-config [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config jwk-config
```

```
show running-config jwk-config [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
one-line
```

DELETE

```
delete jwk-config [name]
```

DESCRIPTION

You can use the jwk-config component to configure a cryptographic JSON Web Key. This key may be used by the Authorization server to sign JSON Web Tokens or by the Client/Resource Server to verify the JSON Web Token signature.

EXAMPLES

```
create jwk-config myjwk {
  alg-type RS256
  key-id b2719f31c6ba1e5fe664fbb1bf0f7c05b3d3a0a1
  modulus ovtSWEWv9Q97JbB5Knfq4iAn8gl-ONzsFoxEasbh9-l4CgeTImIXH31cOxu5tjVjAxeFifPW9w8EdEa-o8kUSJ40Fp2qMRN9wFAHmu5pmS7
  public-exponent AQAB
}
```

Creates a JSON Web Key named myjwk that uses algorithm RS256.

```
create jwk-config myjwk {
  alg-type RS256
  key-id b2719f31c6ba1e5fe664fbb1bf0f7c05b3d3a0a1
  cert myCert.crt
  cert-key myKey.key
}
```

Creates a JSON Web Key named myjwk that will automatically generate other fields based on the values in 'cert' and 'cert-key'. This JSON Web Key uses algorithm RS256 and can be used by the Authorization server to sign JSON Web Tokens.

```
list jwk-config
```

Displays a list of registered JSON Web Keys.

```
delete jwk-config myjwk
```

Deletes the JSON Web Key myjwk

OPTIONS

alg-type

Specifies which cryptographic algorithm is used by this JSON Web Key. The default value is none.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auto-generated

Specifies whether this key was created manually or generated through OpenID Connect metadata discovery. The default value is false.

cert Specifies the certificate this JSON Web Key uses to verify the JWT. Values derived from this field are a part of the JWKS endpoint response.

cert-chain

Specifies the certificate chain this JSON Web Key uses to validate the certificate in the cert field. Values derived from this field are a part of the JWKS endpoint response.

cert-key

Specifies the certificate key this JSON Web Key uses to sign the JWT.

cert-thumbprint-sha1

Specifies the base64url-encoded SHA-1 thumbprint of the DER encoding of X.509 certificate. If the 'cert' field is present, this value is auto-generated.

cert-thumbprint-sha256

Specifies the base64url-encoded SHA-256 thumbprint of the DER encoding of X.509 certificate. If the

'cert' field is present, this value is auto-generated.

curve

Specifies the curve used by the Elliptic Curve JSON Web Key. If the 'cert' field is present, this value is auto-generated.

include-x5c

Specifies whether or not JWKS endpoint response contains a chain of one or more PKIX certificates. The default value is false.

key-id

Specifies the parameter to identify a specific JSON Web Key.

key-type

Specifies the cryptographic algorithm family used by the JSON Web Key. This setting is required. The default value is rsa.

key-use

Specifies whether the JSON Web Key is used for signature generation and verification. At this time, the only supported value is signing.

modulus

Specifies the modulus value for the RSA public key in base64url-encoded format. If the 'cert' field is present, this value is auto-generated.

partition

Displays the partition within which the component resides.

passphrase

Specifies the passphrase used to encrypt the certificate key provided in 'cert-key' field.

public-exponent

Specifies the exponent value for the RSA public key in base64url-encoded format. If the 'cert' field is present, this value is auto-generated.

shared-secret

Specifies the shared secret for the symmetric JSON Web Key when 'key-type' is set to octet.

use-client-secret

Specifies that this JSON Web Key uses client-secret instead of shared-secret. This field is relevant only when key-type is set to octet. The default value is false.

x5c Specifies a chain of one or more PKIX certificates represented as a JSON array of certificate value strings. The JSON array is generated using 'cert' and 'cert-chain' field values.

x-coordinate

Specifies the x coordinate for the Elliptic Curve point in base64url-encoded format. If the 'cert' field is present, this value is auto-generated.

y-coordinate

Specifies the y coordinate for the Elliptic Curve point in base64url-encoded format. If the 'cert' field is present, this value is auto-generated.

SEE ALSO

apm oauth jwt-config

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2016, 2017. All rights reserved.

BIG-IP 2017-10-18 apm oauth jwk-config(1)

apm oauth jwt-config

NAME

jwt-config - Manages JSON web tokens to be used with Client/RS.

MODULE

apm oauth

SYNTAX

Configure the jwt-config component within the oauth module using the following syntax.

CREATE/MODIFY

create jwt-config [name] modify jwt-config [name]

options:

```

access-token-expires-in [integer]
allowed-keys [add | delete | replace-all-with] {
  [name]
}
allowed-signing-algorithms [none | HS256 | HS384 | HS512 | RS256 | RS384 | RS512 | ES256 | ES384]
app-service [[string] | none]
audience [[string] | none]
auto-generated [bool]
blacklist-access-tokens [add | delete | modify | none | replace-all-with] {
  name [string] {
    app-service [[string] | none]
    value-list [add | delete | none | replace-all-with] {
      name [string]
    }
  }
}
blocked-keys [add | delete | replace-all-with] {
  [name]
}
blocked-signing-algorithms [none | HS256 | HS384 | HS512 | RS256 | RS384 | RS512 | ES256 | ES384]
issuer [[string] | none]
jwks-uri [[string] | none]
use-jwt-provider-list-settings [bool]

```

```
edit jwt-config [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```

all-properties
non-default-properties
one-line

```

DISPLAY

```

list jwt-config
list jwt-config [ [ [name] | [glob] | [regex] ] ... ]
show running-config jwt-config
show running-config jwt-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line

```

DELETE

```
delete jwt-config [name]
```

DESCRIPTION

You can use jwt-config component for JWT config management to be used by Client/RS.

EXAMPLES

```

create jwt-config myJwt {
  allowed-keys {
    myJwk1 { }
    myJwk2 { }
    myJwk3 { }
  }
  allowed-signing-algorithms { RS256 }
  issuer https://abc.com
}

```

Creates a JSON web token named myJwt that allows signing algorithm RS256 and JSON web keys myJwk1, myJwk2, myJwk3 and the issuer is https://abc.com.

```
list jwt-config
```

Displays a list of registered JSON web tokens.

```
delete jwt-config myJwt
```

Deletes the JSON web token myJwt.

OPTIONS

access-token-expires-in
Specifies the number of minutes the access token should live. Default value is 0, which means the token never expires.

allowed-keys
Specifies the list of allowed JSON web keys for the token.

allowed-signing-algorithms
Specifies the list of allowed signing algorithms for the token.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

audience
Specifies the audience for the token.

auto-generated

Specifies whether this token was configured manually or was generated through auto-discovery. This is a read-only attribute.

blacklist-access-tokens

Specifies key-value-list that can be used to blacklist tokens based on the key and the list of values for that key.

blocked-keys

Specifies the list of blocked JSON web keys for the token.

blocked-signing-algorithms

Specifies the list of blocked signing algorithms for the token.

issuer

Specifies the issuer of the token.

jwt-uri

Specifies the location of public signing keys for an OAuth Provider. This field is read-only.

use-jwt-provider-list-settings

Specifies whether the settings configured in jwt-provider-list of which this JWT config is a part, should be used. The default value is true.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2016, 2017. All rights reserved.

BIG-IP 2017-06-29 apm oauth jwt-config(1)

apm oauth jwt-provider-list

NAME

jwt-provider-list - Configure a list of providers for JSON web token management to be used by Client/RS.

MODULE

apm oauth

SYNTAX

Configure the jwt-provider-list component within the oauth module using the following syntax.

CREATE/MODIFY

create jwt-provider-list [name] modify jwt-provider-list [name]

options:

access-token-expires-in [integer]

app-service [[string] | none]

providers [add | delete | none | replace-all-with] {

[provider-name]

}

edit jwt-provider-list [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DISPLAY

list jwt-provider-list

list jwt-provider-list [[[name] | [glob] | [regex]] ...]

show running-config jwt-provider-list

show running-config jwt-provider-list [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete jwt-provider-list [name]

DESCRIPTION

You can use the jwt-provider-list component to configure a list of OAuth Providers on Client/RS. A provider list enables a single OAuth Scope agent in an access policy to validate tokens issued by multiple OAuth providers.

EXAMPLES

```
    create jwt-config myProviderList {
  providers {
    myProvider1 { }
    myProvider2 { }
    myProvider3 { }
  }
}
```

Creates a Provider list named myProviderList with Providers myProvider1, myProvider2 and myProvider3.

```
list jwt-provider-list
```

Displays a list of Provider lists.

```
delete jwt-provider-list myProviderList
```

Deletes the Provider List myProviderList.

OPTIONS

access-token-expires-in
Specifies the number of minutes the JSON web token should live. Default value is 0, which means the token never expires.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

providers
Specifies the list of providers that can be used by a single OAuth Scope agent in an access policy to validate tokens issued by multiple OAuth providers.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2016, 2017. All rights reserved.

BIG-IP 2017-06-20 apm oauth jwt-provider-list(1)

apm oauth oauth-claim

NAME

oauth-claim - Manages claims for OAuth Authorization Server.

MODULE

apm oauth

SYNTAX

Configure the oauth-claim component within the oauth module using the following syntax.

CREATE/MODIFY

```
create oauth-claim [name]
modify oauth-claim [name]
options
  app-service [[string] | none]
  claim-description [[string] | none]
  claim-type [boolean | custom | number | string]
  claim-name [string]
  claim-value [[string] | none]
```

```
edit oauth-claim [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
  one-line
```

DISPLAY

```
list oauth-claim
list oauth-claim [ [ [name] | [glob] | [regex] ] ... ]
show running-config oauth-claim
show running-config oauth-claim [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

DELETE
delete oauth-claim [name]

DESCRIPTION

You can use the oauth-claim component to create and manage claims that provide different levels of access control based on end user's role or any other criteria.

EXAMPLES

```
create oauth-claim profileClaim {
  claim-description "Employee Profile"
  claim-type string
  claim-name profile
  claim-value https://company.com/username
}
```

Creates a claim named profileClaim with claim-type set to string, claim-name set to profile and corresponding value being https://company.com/username.

```
list oauth-claim
```

Displays a list of OAuth Claims.

```
delete oauth-claim profileClaim
```

Deletes the OAuth Claim named profileClaim.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

[name]

Specifies the name of the OAuth Claim object. This setting is required.

claim-description

Specifies the description of the claim.

claim-type

Specifies the type of the claim Value.

claim-name

Specifies the name of the claim.

claim-value

Specifies the value of the claim. This value can be any string or session variable.

SEE ALSO

apm policy agent oauth-authz, apm oauth oauth-client-app apm profile oauth

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2018-02-06 apm oauth oauth-claim(1)

apm oauth oauth-client-app

NAME

oauth-client-app - Manages client applications to use with OAuth Authorization Server.

MODULE

apm oauth

SYNTAX

Configure the oauth-client-app component within the oauth module using the following syntax.

CREATE/MODIFY

```
create oauth-client-app [name]
```

```
modify oauth-client-app [name]
```

options:

```
access-token-lifetime [integer]
```

```
app-description [[string] | none]
```

```
app-name [string]
```

```
app-service [[string] | none]
```

```
audience [add | delete | none | replace-all-with] {
```

```

  [string]
}
auth-code-lifetime [integer]
auth-type [none | secret | certificate]
client-cert-dn [[string] | none]
contact [[string] | none]
customization-group [[string] | none]
generate-jwt-refresh-token [true | false]
generate-refresh-token [true | false]
grant-code [enabled | disabled]
grant-password [enabled | disabled]
grant-token [enabled | disabled]
id-token-claims [add | delete | none | replace-all-with] {
  [claim-name]
}
id-token-lifetime [integer]
jwt-access-token-claims [add | delete | none | replace-all-with] {
  [claim-name]
}
jwt-access-token-lifetime [integer]
jwt-refresh-token-lifetime [integer]
logo-url [[string] | none]
openid-connect [enabled | disabled]
redirect-uris [add | delete | none | replace-all-with] {
  [URI]
}
refresh-token-lifetime [integer]
refresh-token-usage-limit [integer]
regenerate-client-secret
reuse-access-token [true | false]
reuse-refresh-token [true | false]
scopes [add | delete | replace-all-with] {
  [scope-name]
}
use-profile-token-mgmt-settings [true | false]
userinfo-claims [add | delete | none | replace-all-with] {
  [claim-name]
}
website-url [[string] | none]
edit oauth-client-app [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line

```

DISPLAY

```

list oauth-client-app
list oauth-client-app [ [ [name] | [glob] | [regex] ] ... ]
show running-config oauth-client-app
show running-config oauth-client-app [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line

```

DELETE

```
delete oauth-client-app [name]
```

DESCRIPTION

You can use the `oauth-client-app` component to register and manage client applications that will make protected resource requests to the OAuth Authorization server on behalf of the resource owner and with its authorization.

EXAMPLES

```

create oauth-client-app myClientApplication {
  app-description "Test App is an application that tests all grant types."
  app-name "Test App"
  grant-code enabled
  grant-password enabled
  grant-token enabled
  logo-url "https://abc.cloud.net/www/public/assets/images/logos/testapp.png"
  redirect-uris add { https://vm1.lab.fp.f5net.com/oauth2/f5_test.php }
  scopes add { scope1 scope2 }
  website-url "https://www.test.com"
  use-profile-token-mgmt-settings false
  audience add { rs1 rs2 }
  jwt-access-token-claims add { claim1 claim2 }
}

```

Creates a client application named `myClientApplication` that will use the generated client credentials to send requests to this Authorization server. It can send token requests using any of the three supported grant types (authorization code, resource owner password credentials or implicit) and uses the default authentication type `"secret"`.

The authorization server will use the configured redirect uri to re-direct back to the client. The client application is associated with configured scopes named `scope1` and `scope2`.

The authorization server will not use the token management settings from the profile, and hence it will use the configured audience rs1 and rs2 and claim claim1 and claim2 values when a JWT access token is returned to the client.

list oauth-client-app

Displays a list of registered client-apps.

delete oauth-client-app myClientApplication

Deletes the OAuth client application myClientApplication

OPTIONS

access-token-lifetime

Specifies the number of minutes for which the access token should be valid. The default is 5 minutes.

app-description

Specifies a user-defined description for the client-app. The default value is none.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

audience

Specifies the audience claim for which the JWT access token is intended. This is a list of values. Each value in this list can be a string, URI, or session variable.

auth-code-lifetime

Specifies the number of minutes for which the authorization code should be valid. The default is 5 minutes.

auth-type

Specifies the authentication type the client will use when it makes requests to the Authorization Server. The default value is secret and other possible values are none and certificate.

client-cert-dn

Specifies the distinguished name of the client certificate that is used to validate a request from client when authentication type is set to certificate.

client-id

Specifies the client ID that uniquely identifies the client application. This field will be auto-generated and should not be specified or modified. Also, this entry cannot be edited once it has been generated.

client-secret

Specifies the client secret that is used to validate a request from client when authentication type is set to secret. This field will be auto-generated and should not be specified or modified.

contact

Specifies a means to contact the developer of the client application.

customization-group

Specifies the customization settings for the client application.

generate-jwt-refresh-token

Specifies whether a refresh token should be generated along with the JWT access token. This is applicable only for "Authorization Code" and "Resource Owner Password Credentials" grant types. The default is true.

generate-refresh-token

Specifies whether a refresh token should be generated along with the access token. This is applicable only for "Authorization Code" and "Resource Owner Password Credentials" grant types.

grant-code

Specifies whether the client application will use the "authorization code" grant type. This grant type must be enabled in order to support hybrid flow in OpenID Connect. The default value is disabled. At least one grant type must be set to enabled.

grant-password

Specifies whether the client application will use the "resource owner password credentials" grant type. The default value is disabled.

grant-token

Specifies whether the client application will use the "implicit" grant type. The default value is disabled.

id-token-claims

Specifies the list of claims that are part of ID token.

id-token-lifetime

Specifies the number of minutes for which the ID token should be valid. The default is 5 minutes.

jwt-access-token-claims

Specifies the list of claims that are part of JWT access token.

jwt-access-token-lifetime

Specifies the number of minutes for which the JWT access token should be valid. The default is 5 minutes.

`jwt-refresh-token-lifetime`

Specifies the number of minutes for which the JWT refresh token should be valid. The default is 60 minutes.

`logo-url`

Specifies the path from which the logo of the client application can be displayed.

`openid-connect`

Specifies whether this client app supports OpenID Connect or not.

`[name]`

Specifies the name of the OAuth Client Application. This setting is required.

`partition`

Displays the partition within which the component resides.

`redirect-uris`

Specifies the list of re-direct URIs that the Authorization Server will use to re-direct back to the client after processing a request. This setting should have at least one entry if the client application uses the authorization code grant type or the implicit grant type.

`refresh-token-lifetime`

Specifies the number of minutes for which the refresh token should be valid. The default is 480 minutes.

`refresh-token-usage-limit`

Specifies the maximum number of times the access token can be obtained using the refresh token request. The default value is 64. Value 0 represents unlimited number of times.

`regenerate-client-secret`

Indicates a request to regenerate the client secret. Do not use other means to modify the secret.

`reuse-access-token`

Specifies whether an access token is reused or a new access token is generated when it is obtained using refresh token request. When the access token is reused, its expiry time is extended.

`reuse-refresh-token`

Specifies whether a refresh token is reused or a new refresh token is generated when it is obtained using refresh token request.

`scopes`

Specifies the list of scopes that is to be associated with the client application.

`use-profile-token-mgmt-settings`

Specifies whether the default settings that come from OAuth profile must be used or not.

`userinfo-claims`

Specifies the list of claims that are part of UserInfo.

`website-url`

Specifies the website URL of the client application.

SEE ALSO

`apm profile oauth`

`apm oauth oauth-scope`

`apm oauth oauth-claim`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2017. All rights reserved.

BIG-IP 2017-10-31 `apm oauth oauth-client-app(1)`

apm oauth oauth-resource-server

NAME

`oauth-resource-server` - Manages resource servers to use with OAuth Authorization Server.

MODULE

`apm oauth`

SYNTAX

Configure the `oauth-resource-server` component within the `oauth` module using the following syntax.

CREATE/MODIFY

create oauth-resource-server [name]

modify oauth-resource-server [name]

options:

app-service [[string] | none]

description [[string] | none]

auth-type [none | secret | certificate]

regenerate-resource-server-secret

resource-server-cert-dn [[string] | none]

edit oauth-resource-server [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DISPLAY

list oauth-resource-server

list oauth-resource-server [[name] | [glob] | [regex]] ...]

show running-config oauth-resource-server

show running-config oauth-resource-server [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete oauth-resource-server [name]

DESCRIPTION

You can use the oauth-resource-server component to register and manage resource servers that host resources that will be accessed by the user. Resource servers can accept and respond to protected resource requests using access tokens.

EXAMPLES

```
create oauth-resource-server myResourceServer {
resource-server-cert-dn "/C=US/ST=CA/L=SJ/O=Company Name,Inc/OU=Engg/CN=user-name/emailAddress=username@company-domain."
}
```

Creates a resource server named myResourceServer that will use the generated resource server credentials to send requests to the Authorization server. It uses the default authentication type "certificate".

```
list oauth-resource-server
```

Displays a list of registered resource servers.

```
delete oauth-resource-server myResourceServer
```

Deletes the OAuth resource server myResourceServer.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auth-type

Specifies the authentication type the resource server will use when it makes requests to the Authorization Server. The default value is certificate and other possible values are none and secret.

[name]

Specifies the name of the OAuth Resource Server. This setting is required.

description

Specifies the description of the OAuth Resource Server Object.

partition

Displays the partition within which the component resides.

regenerate-resource-server-secret

Indicates a request to regenerate the resource server secret. Do not use other means to modify the secret.

resource-server-cert-dn

Specifies the distinguished name of the resource server certificate that is used to validate a request from the resource server when authentication type is set to certificate.

resource-server-id

Specifies the resource server ID that uniquely identifies the resource server. This field will be auto-generated and should not be specified or modified. Also, this entry cannot be edited after it has been generated.

resource-server-secret

Specifies the resource server secret that is used to validate a request from the resource server when authentication type is set to secret. This field will be auto-generated and should not be specified or modified.

SEE ALSO

apm profile oauth

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2016. All rights reserved.

BIG-IP 2016-11-08 apm oauth oauth-resource-server(1)

apm oauth oauth-scope

NAME

oauth-scope - Manages scopes for OAuth Authorization Server.

MODULE

apm oauth

SYNTAX

Configure the oauth-scope component within the oauth module using the following syntax.

CREATE/MODIFY

```
create oauth-scope [name]
```

```
modify oauth-scope [name]
```

options

```
app-service [[string] | none]
```

```
customization-group [[string] | none]
```

```
scope-description [[string] | none]
```

```
scope-name [string]
```

```
scope-value [[string] | none]
```

```
edit oauth-scope [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DISPLAY

```
list oauth-scope
```

```
list oauth-scope [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config oauth-scope
```

```
show running-config oauth-scope [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DELETE

```
delete oauth-scope [name]
```

DESCRIPTION

You can use the oauth-scope component to create and manage scopes that provide different levels of access control based on end user's role or any other criteria.

EXAMPLES

```
create oauth-scope myOAuthScope {
  customization-group "company_my_oauth_scope"
  scope-description "Group in company"
  scope-name "profile"
  scope-value "Product Development"
}
```

Creates a scope named myOAuthScope that uses customization group company_my_oauth_scope to customize scope information in the OAuth Authorization page. It is configured with profile and Product Development as its name and value respectively.

```
list oauth-scope
```

Displays a list of OAuth Scopes.

```
delete oauth-scope myOAuthScope
```

Deletes the OAuth Scope named myOAuthScope.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot

modify or delete the object. Only the application service can modify or delete the object.

`customization-group`
Specifies the customization settings for the scope.

`[name]`
Specifies the name of the OAuth Scope. This setting is required.

`partition`
Displays the partition within which the component resides.

`scope-description`
Specifies the description of the scope.

`scope-name`
Specifies the name of the scope.

`scope-value`
Specifies the value of the scope. This value can be any string or session variable.

SEE ALSO

`apm policy agent oauth-authz`, `apm oauth oauth-client-app`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2016. All rights reserved.

BIG-IP 2016-06-30 `apm oauth oauth-scope(1)`

apm oauth purged-entries

NAME

Purged-Entries - Displays OAuth records purged from the respective database instance.

MODULE

`apm oauth purged-entries`

SYNTAX

Displays all the purged-entries within the oauth module using the syntax shown below.

DISPLAY

```
show purged-entries
show purged-entries [db-instance]
options:
  all-properties
  non-default-properties
  one-line
  recursive
output only properties:
  oauth_id [string]
  purged_time [string]
```

DESCRIPTION

You can use the purged-entries component to manage purged records.

EXAMPLES

```
show apm oauth purged-entries
```

Lists all the OAuth purged entries for the oauth default db instance `oauthdb`. Also displays the total number of records purged.

```
show apm oauth purged-entries db-instance test
```

Lists all the purged entries for the oauth db instance `test`. Also displays the total number of records purged.

OPTIONS

`db-instance`

This specifies the db instance. If no value is specified, the default db instance, `oauthdb`, is used.

`oauth-id`

Specifies oauth id. Displays in the output of the show command.

`purged_time`

Specifies the timestamp at which each record was purged. The format is "YYYY-MM-DD HH:MM:SS". Displays in the output of the show command.

SEE ALSO

apm oauth db-instance

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2016-11-07 apm oauth purged-entries(1)

apm oauth token-details

NAME

Token-Details - Displays OAuth token properties and revokes access tokens and refresh tokens.

MODULE

apm oauth token-details

SYNTAX

Lists the token-details within the oauth module using the syntax shown in the following sections.

DISPLAY

list token-details

list token-details [[[oauth-id] | [app-name] | [client-id] | [offset] | [db-instance]] ...]

options:

all-properties
non-default-properties
one-line
recursive

output only properties:

access-token-status [[string] | none]
access-token-issued-at [[string] | none]
access-token-expires-at [[string] | none]
refresh-token-status [[string] | none]
refresh-token-issued-at [[string] | none]
refresh-token-expires-at [[string] | none]
grant-type [[string] | none]
user-agent [[string] | none]

show token-details

show token-details [[[app-name] | [client-id] | [offset] | [db-instance]] ...]

options:

field-fmt
recursive

REVOKE

revoke token-details [[oauth-id client-id]]

DESCRIPTION

You can use the token-details component to manage oauth tokens.

EXAMPLES

list apm oauth token-details

Lists all the token details for the oauth default db instance of oauthdb.

list apm oauth token-details db-instance test

Lists all the token details for the oauth db instance test.

list apm oauth token-details client-app myapp

Lists all the token details for the client app myapp in the default db instance.

show apm oauth token-details

Shows the number of tokens in the oauth default db instance of oauthdb.

show apm oauth token-details db-instance test

Shows the number of tokens in oauth db instance test.

show apm oauth token-details client-app myapp

Shows the number of token for the client app myapp in default db instance.

revoke apm oauth token-details 93b225478484f0fa0addc527b2c50001d7e3yyb70c3eda99e56 client-id
ae8b6cd3708a5805e84ae48369ab0001d7eb70c1325c8d56

Revokes the access/refresh token for oauth id 93b225478484f0fa0addc527b2c50001d7eb70c3eda99e56 and client id ae8b6cd3708a5805e84ae48369ab0001d7eb70c1325c8d56.

revoke apm oauth token-details 93b225478484f0fa0addc527b2c50001d7e3yyb70c3eda99e56 db-instance test client-id
ae8b6cd3708a5805e84ae48369ab0001d7eb70c1325c8d56

Revokes the access/refresh token for oauth id 93b225478484f0fa0addc527b2c50001d7eb70c3eda99e56 and client

id ae8b6cd3708a5805e84ae48369ab0001d7eb70c1325c8d56 in the db instance test.

OPTIONS

client-id

Specifies the client id for which the token is generated. Must be specified for revoke command.

app-name

Specifies the app name for which the token is generated. This is an optional parameter for list and show commands.

offset

Specifies the number of tokens to skip before beginning to display the remaining tokens.

Example: list apm oauth token-details offset 10

This will return the tokens from the count of 11 and will skip the first 10.

db-instance

This specifies the db instance. If no value is specified, oauthdb (the default db instance) is used.

user-id

Specifies the user id for which the token is generated.

oauth-id

Specifies oauth id. Displays in the output of the list command.

access-token-status

Specifies the access token status, such as active,inactive,expired and revoked. Displays in the output of the list command.

access-token-issued-at

Specifies the access token issued timestamp. The format is "YYYY-MM-DD HH:MM:SS". Displays in the output of the list command.

access-token-expires-at

Specifies the access token expiry timestamp. The format is "YYYY-MM-DD HH:MM:SS". Displays in the output of the list command.

refresh-token-status

Specifies the refresh token status such as active,inactive,expired and revoked. Displays in the output of the list command.

refresh-token-issued-at

Specifies the refresh token issued timestamp. The format is "YYYY-MM-DD HH:MM:SS". Displays in the output of the list command.

refresh-token-expires-at

Specifies the refresh token expiry timestamp. The format is "YYYY-MM-DD HH:MM:SS". Displays in the output of the list command.

grant-type

Specifies the grant type; valid values are ropc, implicit and auth-code. Displays in the output of the list command.

user-agent

Specifies the user agent information for the token. Displays in the output of the list command.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2017-04-25 apm oauth token-details(1)

apm policy access-policy

NAME

access-policy - Manages an access policy.

MODULE

apm policy

SYNTAX

Warning: F5 Networks recommends that you use the visual policy editor in the Configuration utility to create and manage access policies. If you are using tmsh scripts to create an access policy, please do the following:

1. Create transaction.
2. Create policy agents; each policy must include an ending agent type.
3. Create access policy items for start and end and attach the corresponding agents.
4. Create access policy and add all the policy items.
5. Create access profile with the corresponding access policy.
6. Submit transaction.

Failure to follow the above mentioned steps will result in error messages or invalid configuration. This sample script creates a policy with start and end agents only:

1. create cli transaction
2. create apm policy agent ending-allow /Common/rest_end_allow_ag { }
- 3a. create apm policy policy-item /Common/rest_end_allow { agents add { /Common/rest_end_allow_ag { type ending-allow } } caption Allow
- 3b. create apm policy policy-item /Common/rest_ent { caption Start color 1 rules { { caption fallback next-item /Common/rest_end_allow } }
4. create apm policy access-policy /Common/rest { default-ending /Common/rest_end_allow items add { rest_end_allow { } rest_ent { } } sta
5. create apm profile access /Common/rest { accept-languages add { en } access-policy /Common/rest log-settings add { default-log-setting
6. submit cli transaction

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2016-12-01 apm policy access-policy(1)

apm policy agent aaa-active-directory

NAME

aaa-active-directory - Manages an AAA Active Directory(r) agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-active-directory component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create aaa-active-directory [name]
modify aaa-active-directory [name]
options
  app-service [[string] | none]
  auth-max-logon-attempt [integer]
  fetch-nested-groups [true | false]
  fetch-primary-groups [true | false]
  hints [true | false]
  query-attrname [[string] | none]
  query-filter [[string] | none]
  server [[string] | none]
  trusted-domains [[string | none]]
  show-extended-error [true | false]
  type [query | auth | last]
  upn [true | false]
```

DISPLAY

```
list aaa-ldap
list aaa-ldap [ [ [name] | [glob] | [regex] ] ... ]
show running-config aaa-ldap
show running-config aaa-ldap [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  current-module
  non-default-properties
  one-line
  app-service
  partition
```

DELETE

```
delete aaa-active-directory ([name] | all)
```

DESCRIPTION

You can use the aaa-active-directory component to configure an AAA Active Directory agent.

EXAMPLES

```
create aaa-active-directory MyADQueryagent {query-filter "(be sAMAccountName=%{session.logon.last.username})"
type query server "companyAD" }
```

Creates the query type AAA Active Directory agent named MyADQueryagent that uses the (be SAMAccountName=%{session.logon.last.username}) filter and the companyAD AAA AD Server.

create agent aaa active MyADAuthagent { type auth server "companyAD" }
Creates the authorization type AAA Active Directory agent named MyADAuthagent that uses the companyAD AAA AD server.

list aaa-active-directory all
Displays a list of AAA Active Directory agents and their properties.

delete aaa-active-directory MyADagent
Deletes the MyADagent AAA Active Directory agent.

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auth-max-logon-attempt
Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

fetch-nested-groups
When enabled, the system administrator can retrieve the full list of groups that user belongs to, even if the retrieval privileges are nested through other groups to which the user belongs to directly. The default value is false.

fetch-primary-groups
When enabled, the system administrator can retrieve the primary group of a user, and use that name as a group in access policy item rules. The default value is false.

hints
When enabled, the system offers the user an option to create a hint that assists in remembering a password. The default value is false.

query-attrname
Specifies the attribute name that you are adding or deleting for the agent.

query-filter
Specifies the search criteria the system uses when querying an AAA Active Directory(r) server for authentication information. The system supports session variables as part of search query string.

[name]
Specifies the name of an AAA Active Directory agent. This setting is required.

partition
Displays the partition within which the component resides.

server
Specifies an AAA Active Directory server the system uses for Active Directory queries and authentication.

server
Specifies an AAA Active Directory Trusted Domains object that the system uses for Active Directory queries and authentication. This option requires upn option to be enabled

show-extended-error
Specifies to display a verbose error message. The default value is false.

type Specifies the type of AAA Active Directory agent. The default value is last.

query
Specifies that the agent makes a query against the AAA Active Directory Server to retrieve information in accordance with the query-filter and query-attributes options.

auth Specifies that the agent is an authentication agent only. It uses the AAA Active Directory Server, but only for authentication purposes. APM does not get any information from the Domain.

last
upn When enabled, APM supports the user principal name (UPN) naming style and process cross-domain authentication requests. Some examples of UPNs are: user@fqdn.of.domain.com, user@upnsuffix.com, and user@domain. The default value is false.

SEE ALSO
tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2013-11-15 apm policy agent aaa-active-directory(1)

apm policy agent aaa-client-cert

NAME

aaa-client-cert - Manages an AAA Client Certification agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-client-cert component within the policy-agent module using the following syntax.

CREATE/MODIFY

modify aaa-client-cert [name]

create aaa-client-cert [name]

options:

app-service [[string] | none]

mode [request | require]

DISPLAY

list aaa-client-cert

list aaa-client-cert [[[name] | [glob] | [regex]] ...]

show running-config aaa-client-cert

show running-config aaa-client-cert [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

current-module

non-default-properties

one-line

app-service

partition

DELETE

delete aaa-client-cert [name]

DESCRIPTION

You can use this component to configure an AAA Client Certification agent.

EXAMPLES

create aaa-client-cert MyCCagent

Creates the AAA Client Certification agent named MyCCagent in the Common partition.

list aaa-client-cert all

Displays a list of AAA Client Certification agents.

delete aaa-client-cert MyCCagent

Deletes the MyCCagent AAA Client Certification agent.

OPTIONS

[name]

Specifies the name of an AAA client cert agent. This setting is required.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

mode Specifies the mode (request/require) for this certificate. The options are:

request

Specifies that the system requests a valid certificate from a client, but always authenticates the client.

require

Specifies that the system requires a client to present a valid certificate.

partition

Displays the partition within which the component resides.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

apm policy agent aaa-crldp

NAME

aaa-crldp - Manages an AAA CRLDP (Constraint-Based Routed Label Distributed Protocol) agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-crldp component within the policy agent module using the following syntax.

CREATE/MODIFY

create aaa-crldp [name]

modify aaa-crldp [name]

options:

app-service [[string] | none]

server (| none)

DISPLAY

list aaa-crldp

list aaa-crldp [[[name] | [glob] | [regex]] ...]

show running-config aaa-crldp

show running-config aaa-crldp [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

app-service

current-module

non-default-properties

one-line

partition

DELETE

delete aaa-crldp [name]

DESCRIPTION

You can use the aaa-crldp component to create and manage an AAA CRLDP agent.

EXAMPLES

create aaa-crldp MyCCagent

Creates an AAA CRLDP agent named MyCCagent in the Common partition.

list aaa-crldp all

Displays a list of AAA CRLDP agents.

delete aaa-crldp MyCCagent

Deletes the MyCCagent AAA CRLDP agent.

OPTIONS

[name]

Specifies the name of an agent that you want to display or delete. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

server

Specifies the name of the server on which this agent resides. This option is required.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

apm policy agent aaa-http

NAME

aaa-http - Manages an AAA HTTP agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-http component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create aaa-http [name]
modify aaa-http [name]
options
  app-service [[string] | none]
  max-logon-attempt [integer]
  server [[string] | none]
```

DISPLAY

```
list aaa-http
list aaa-http [ [ [name] | [glob] | [regex] ] ... ]
show running-config aaa-http
show running-config aaa-http [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  current-module
  non-default-properties
  one-line
  app-service
  partition
```

DELETE

```
delete aaa-http [name]
```

DESCRIPTION

You can use the aaa-http component to configure an AAA HTTP agent.

EXAMPLES

```
create aaa-http MyCCagent
Creates the aaa-http agent named MyCCagent in the Common partition.
```

```
list all aaa-http
Displays a list of aaa-http agents.
```

```
delete aaa-http MyCCagent
Deletes the MyCCagent aaa-http agent.
```

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

max-logon-attempt
Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

[name]
Specifies the name of an AAA HTTP agent. This setting is required.

partition
Displays the partition within which the component resides.

server
Specifies which AAA HTTP server the system uses for Active Directory queries and authentication.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

apm policy agent aaa-ldap

NAME

aaa-ldap - Manages an AAA LDAP(r) agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-ldap component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create aaa-ldap [name]
modify aaa-ldap [name]
options:
  app-service [[string] | none]
  attr-name ( | none) [add | delete]
  filter [[string] | none]
  group-member-scope [none | direct | all]
  group-membership-scope [none | direct | all]
  max-logon-attempt [integer]
  search-dn [[string] | none]
  server [[string] | none]
  show-extended-error [true | false]
  type [query | auth | modify | last]
  user-dn [[string] | none]
  modify-type [add | modify | delete | modify-last]
  ldapmod-attributes ( | none) [add | delete]
```

DISPLAY

```
list aaa-ldap
list aaa-ldap [ [name] | [glob] | [regex] ] ... ]
show running-config aaa-ldap
show running-config aaa-ldap [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  current-module
  non-default-properties
  one-line
  app-service
  partition
```

DELETE

```
delete aaa-ldap [name]
```

DESCRIPTION

Use this component to create, modify, display, or delete an AAA LDAP agent.

EXAMPLES

```
create aaa-ldap MyLDAPagent { user-dn "cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=com" type
auth server "companyLDAP" } aaa-ldap MyLDAPagent { search-dn "cn=users,dc=lab,dc=fp,dc=com" filter
"(SAMAccountName=%{{session.logon.last.username}})" type auth server "companyLDAP" }
Creates the authorization type AAA LDAP agent named MyLDAPagent that is associated with the companyLDAP
server that uses the cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=f5net,dc=com user domain
name, the cn=users,dc=lab,dc=fp,dc=com search domain, and the
(SAMAccountName=%{{session.logon.last.username}}) filter.
```

```
create aaa-ldap MyLDAPagent { search-dn "cn=users,dc=lab,dc=fp,dc=com" filter
"(sAMAccountName=%{{session.logon.last.username}})" type query server "companyLDAP" }
Creates the query type AAA LDAP agent named MyLDAPagent that is associated with the companyLDAP server
that uses the cn=users,dc=lab,dc=fp,dc=com search domain and the
(SAMAccountName=%{{session.logon.last.username}}) filter.
```

```
create aaa-ldap MyLDAPagent { user-dn "cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=com" type
modify modify-type add server "companyLDAP" ldapmod-attributes add { objectClass { mod-op add mod-values add {
top person organizationalPerson user } } cn { mod-op add mod-values add { demo } } } }
Creates the modify type AAA LDAP agent named MyLDAPagent that is associated with the companyLDAP server
that uses the cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=f5net,dc=com user domain name,
the add modify type, and the ldapmod attributes
```

```
create aaa-ldap MyLDAPagent { user-dn "cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=com" type
modify modify-type modify server "companyLDAP" ldapmod-attributes add { givenName { mod-op replace mod-values
add { demo } } } }
Creates the modify type AAA LDAP agent named MyLDAPagent that is associated with the companyLDAP server
that uses the cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=f5net,dc=com user domain name,
the modify modify type, and the ldapmod attributes which uses givenName modify attribute replace mod
operation and the demo mod values
```

```
create aaa-ldap MyLDAPagent { user-dn "cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=com" type
modify modify-type delete server "companyLDAP" }
```

Creates the modify type AAA LDAP agent named MyLDAPagent that is associated with the companyLDAP server that uses the cn=%{session.logon.last.username},cn=users,dc=lab,dc=fp,dc=f5net,dc=com user domain name, the delete modify type

```
list aaa-ldap
```

Displays a list of AAA LDAP agents.

```
delete aaa-ldap MyLDAPagent
```

Deletes the MyLDAPagent AAA LDAP agent.

OPTIONS

```
app-service
```

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

```
attr-name
```

Adds an attribute name to the agent or deletes an attribute name from the agent.

```
group-member-scope
```

Specifies the scope of user lookup for a group. When the search returns a group, this attribute specifies whether to also look up the members of the group. The options are:

none No members required.

direct Only direct members required.

all All members required. This includes those that derive membership in this group through membership in other groups and those that are direct members.

```
group-membership-scope
```

Specifies the scope of group lookup for a user or a group. When the search returns a user or a group, this attribute specifies whether to also look up the groups to which this user or group belong. The options are:

none No groups required.

direct Only the groups to which the current user or group belong directly are required.

all All groups required. This includes the groups to which the user or the group belong directly and the groups to which the user or group belong indirectly (through membership in another group).

```
filter
```

Specifies the LDAP filter that APM uses when querying an AAA LDAP server for authentication information. You must use the filter option with the search-dn option.

```
max-logon-attempt
```

Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

```
[name]
```

Specifies the name of an AAA LDAP agent. This setting is required.

```
partition
```

Displays the partition within which the component resides.

```
search-dn
```

Specifies the base domain name that APM uses for internal LDAP search operations. You must use the search-dn option with the filter option.

```
server
```

Specifies the AAA LDAP server that the system uses for LDAP queries and authentication.

```
show-extended-error
```

Specifies to display a verbose error message. The default value is false.

```
type
```

Specifies a type of AAA LDAP agent. This setting is required. The default is last.

```
user-dn
```

Specifies the fully qualified domain name of the Access Policy Manager. F5 Networks recommends that you specify this value in lower case and without spaces for compatibility with some specific LDAP servers. The specific content of this string depends on your directory layout.

SEE ALSO

tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 apm policy agent aaa-ldap(1)

apm policy agent aaa-oauth

NAME

aaa-oauth - Manages an AAA OAuth(r) agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-oauth component within the policy agent module using the following syntax.

CREATE/MODIFY

create aaa-oauth [name]

modify aaa-oauth [name]

options:

app-service [[string] | none]

auth-redirect-request [name]

grant-type [authorization-code | password]

redirection-uri [string]

response [name]

scope [[string] | none]

scope-data-request [name]

server [name]

token-refresh-request [name]

token-request [name]

type [client | scope]

validation-scopes-request [name]

DISPLAY

list aaa-oauth

list aaa-oauth [[[name] | [glob] | [regex]] ...]

show running-config aaa-oauth

show running-config aaa-oauth [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

current-module

non-default-properties

one-line

app-service

partition

DELETE

delete aaa-oauth [name]

DESCRIPTION

Use this component to create, modify, display, or delete an OAuth Client or OAuth Scope agent.

EXAMPLES

```
create aaa-oauth MyGoogleClient { auth-redirect-request GoogleAuthRedirectRequest grant-type authorization-code scope "https://www.googleapis.com/auth/userinfo.email https://www.googleapis.com/auth/userinfo.profile" server myGoogleServer token-request GoogleTokenRequest type client validation-scopes-request GoogleValidationScopesRequest } Creates the OAuth Client agent to acquire an access_token from Google authorization server using authorization-code grant type. Defines two scopes. The user's permission will be requested for the scopes.
```

```
create aaa-oauth MyGoogleScope { scope-data-request { https://www.googleapis.com/auth/userinfo.profile { request GoogleScopeUserInfoProfileRequest } } server myGoogleServer type scope validation-scopes-request GoogleValidationScopesRequest }
```

Creates OAuth Scope agent to get the list of scopes associated with the access_token, and defines the scope-data-request to retrieve more information about user identity if the access_token contains the scope "https://www.googleapis.com/auth/userinfo.profile".

```
list aaa-oauth
```

Displays a list of OAuth agents.

```
delete aaa-oauth MyGoogleScope
```

Deletes the MyGoogleScope OAuth agent.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auth-redirect-request

OAuth Request name to be used for user redirect in order to obtain authorization code.

grant-type

Specifies grant type that should be used to request an access_token.

redirection-uri

Specifies redirection URI. The redirection URI is used by the Authorization Server to redirect user back after authentication. The URI is a property of client application registered at authorization server.

This option is used along with 'authorization-code' grant type only.

response
Specifies the response config object name.

scope
The list of scopes to request user's permission for.

scope-data-request
Defines OAuth Request to obtain additional information from the resource server for the specified scope, using `access_token`.

server
Specifies OAuth Server that represents the authorization server to work with.

token-refresh-request
Specifies OAuth Request to refresh an expired `access_token`.

token-request
Specifies OAuth Request to request an `access_token`.

type Type of the OAuth agent. Available options are: `client` or `scope`. Default value `client`. The type cannot be changed for an existing OAuth agent.

validation-scopes-request
Specifies OAuth Request to validate the `access_token` (when agent type is `client`) or to retrieve list of scopes associated with the `access_token` (when agent type is `scope`).

SEE ALSO
`tmsb`

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2014, 2016. All rights reserved.

BIG-IP 2018-07-12 `apm policy agent aaa-oauth(1)`

apm policy agent aaa-ocsp

NAME
`aaa-ocsp` - Manages an AAA OCSP (Online Certificate Status Protocol) agent.

MODULE
`apm policy agent`

SYNTAX
Configure the `aaa-ocsp` component within the `policy agent` module using the following syntax.

CREATE/MODIFY
`create aaa-ocsp [name]`
`modify aaa-ocsp [name]`
options:
`app-service [[string] | none]`
`certificate-type`
`ocsp-responder`

DISPLAY
`list aaa-ocsp`
`list aaa-ocsp [[[name] | [glob] | [regex]] ...]`
`show running-config aaa-ocsp`
`show running-config aaa-ocsp [[[name] | [glob] | [regex]] ...]`
options:
`all`
`all-properties`
`current-module`
`non-default-properties`
`one-line`
`app-service`
`partition`

DELETE
`delete aaa-ocsp [name]`

DESCRIPTION
Use this command to create, modify, display, or delete an AAA OCSP agent.

EXAMPLES

```
create aaa-ocsp MyCCagent
Creates the AAA OCSF agent named MyCCagent in the Common partition.
```

```
list aaa-ocsp all
Displays a list of AAA OCSF agents.
```

```
delete aaa-ocsp MyCCagent
Deletes the MyCCagent AAA OCSF agent.
```

OPTIONS

[name]
Specifies the name of an agent that you want to display or delete. This setting is required.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

certificate-type
Specifies the type of certificate to check against OCSF responder. The value can be either user or machine. The default value is user.

ocsp-responder
Specifies which OCSF responder object to use to validate a certificate.

partition
Displays the partition within which the object resides.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2015-05-12 apm policy agent aaa-ocsp(1)

apm policy agent aaa-radius

NAME

aaa-radius - Manages an AAA RADIUS agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-radius component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create aaa-radius [name]
modify aaa-radius [name]
options:
  app-service [[string] | none]
  max-logon-attempt
  password-source [[string] | none]
  server ( | none)
  show-extended-error (true | false)
  username-source [[string] | none]
```

DISPLAY

```
list aaa-radius
list aaa-radius [ [name] | [glob] | [regex] ] ... ]
show running-config aaa-radius
show running-config aaa-radius [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  current-module
  non-default-properties
  one-line
  app-service
  partition
```

DELETE

```
delete aaa-radius [name]
```

DESCRIPTION

Use this command to create, modify, display, or delete an AAA RADIUS agent.

EXAMPLES

```
create aaa-radius Myradiusagent {server "companyradius"}
```

Creates an AAA RADIUS agent named Myradiusagent that is associated with the companyradius server.

```
list aaa-radius
```

Displays a list of AAA RADIUS agents.

```
delete aaa-radius Myradiusagent
```

Deletes the Myradiusagent AAA RADIUS agent.

OPTIONS

```
app-service
```

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

```
max-logon-attempt
```

Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

```
[name]
```

Specifies the name of an AAA RADIUS agent. This setting is required.

```
partition
```

Displays the partition within which the object resides.

```
password-source
```

Specifies the session variable name from which RADIUS agent should read the password. The default value is `%{session.logon.last.password}`.

```
server
```

Specifies the AAA RADIUS server that the system uses for RADIUS queries and authentication.

```
show-extended-error
```

Specifies to display a verbose error message. The default value is false.

```
username-source
```

Specifies the session variable name from which RADIUS agent should read the username. The default value is `%{session.logon.last.username}`.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2017-02-08 apm policy agent aaa-radius(1)

apm policy agent aaa-saml

NAME

aaa-saml - Manages a AAA SAML agent.

MODULE

apm policy agent

SYNTAX

Configure the aaa-saml component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create aaa-saml [name]
```

```
modify aaa-saml [name]
```

options:

```
app-service [[string] | none]
```

```
attr-consuming-service-session-var ( | none)
```

```
attribute-consuming-service ( | none)
```

```
server ( | none)
```

DISPLAY

```
list aaa-saml
```

```
list aaa-saml [ [ [name] | [glob] | [regex] ] ... ]
show running-config aaa-saml
show running-config aaa-saml [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  current-module
  non-default-properties
  one-line
  app-service
  partition
```

```
DELETE
delete aaa-saml [name]
```

DESCRIPTION
Use this command to create, modify, display, or delete a AAA SAML agent.

EXAMPLES
create aaa-saml Mysamlagent {server "companysaml"}
Creates a AAA SAML agent named Mysamlagent that is associated with the companysaml server.

```
list aaa-saml
Displays a list of AAA SAML agents.
```

```
delete aaa-saml Mysamlagent
Deletes the Mysamlagent AAA SAML agent.
```

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

```
attr-consuming-service-session-var
Specifies the session variable used as the source for the attribute consuming service index. The index will be included in the SAML authentication request generated. The IdP maps the index to the list of attributes derived from the metadata previously shared and returns those attributes in the SAML Response. attr-consuming-service-session-var and attribute-consuming-service cannot be configured at the same time.
```

```
attribute-consuming-service
Specifies the name of one of the attribute consuming services associated with the server. The index associated with the selected attribute consuming service will be included in the SAML authentication request generated. The IdP maps the index to the list of attributes derived from the metadata previously shared and returns those attributes in the SAML Response. attribute-consuming-service and attr-consuming-service-session-var cannot be configured at the same time.
```

```
partition
Displays the partition within which the object resides.
```

```
server
Specifies the AAA SAML server that the system uses for SAML queries and authentication.
```

SEE ALSO
tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2017-02-03 apm policy agent aaa-saml(1)

apm policy agent aaa-securid

NAME
aaa-securid - Manages an AAA SecurID agent.

MODULE
apm policy agent

SYNTAX
Configure the aaa-securid component within the policy agent module using the following syntax.

```
CREATE/MODIFY
create aaa-securid [name]
modify aaa-securid [name]
```

options:
app-service [[string] | none]
max-logon-attempt [integer]
password-source [[string] | none]
server [[string] | none]
show-extended-error [true | false]
username-source [[string] | none]

edit aaa-securid [[glob] | [regex]] ...]

options:
all-properties
non-default-properties

DISPLAY

list aaa-securid

list aaa-securid [[[name] | [glob] | [regex]] ...]

show running-config aaa-securid

show running-config aaa-securid [[[name] | [glob] | [regex]] ...]

options:
all
all-properties
current-module
non-default-properties
one-line
app-service
partition

DELETE

delete aaa-securid [name]

DESCRIPTION

You can use the aaa-securid component to create and manage an AAA SecurID agent.

EXAMPLES

```
create aaa-securid mySecuridAgent { server rsa1_106 }
```

Creates an AAA SecurID agent named mySecuridAgent that is associated to AAA RSA Server rsa1_106.

```
list all aaa-securid
```

Displays a list of AAA SecurID agents.

```
delete aaa-securid MyCCagent
```

Deletes the MyCCagent AAA Client Certification agent.

OPTIONS

[name]

Specifies the name of an agent that you want to display or delete. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

max-logon-attempt

Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

password-source

Specifies the session variable name from which RSA SecurID agent should read the password. The default value is %session.logon.last.password}.

server

Specifies the AAA RSA SecurID server that the system uses for LDAP queries and authentication.

show-extended-error

Specifies to display a verbose error message. The default value is false.

username-source

Specifies the session variable name from which RSA SecurID agent should read the username. The default value is %session.logon.last.username}.

SEE ALSO

tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

apm policy agent acct-radius

NAME

acct-radius - Manages a RADIUS Accounting agent.

MODULE

apm policy agent

SYNTAX

Configure the acct-radius component within the policy agent module using the following syntax.

CREATE/MODIFY

create acct-radius [name]

modify acct-radius [name]

options:

app-service [[string] | none]

server [[string] | none]

username-source [[string] | none]

edit acct-radius [[glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list acct-radius

list acct-radius [[[name] | [glob] | [regex]] ...]

show running-config acct-radius

show running-config acct-radius [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

current-module

non-default-properties

one-line

app-service

partition

DELETE

delete acct-radius [name]

DESCRIPTION

You can use the acct-radius component to create and manage a RADIUS Accounting agent.

EXAMPLES

```
create acct-radius MyRADIUSagent { server "MyRADIUS" }
```

Creates the MyRADIUSagent RADIUS Accounting agent that is associated with the MyRADIUS server.

```
list acct-radius
```

Displays a list of RADIUS Accounting agents and the servers associated with the agents.

```
delete acct-radius MyRADIUSagent
```

Deletes the MyRADIUSagent RADIUS Accounting agent.

OPTIONS

[name]

Specifies the name of a RADIUS Accounting server. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

server

Specifies a RADIUS Accounting server the system uses to send RADIUS Accounting START and RADIUS Accounting STOP messages. This option is required.

username-source

Specifies the session variable name from which RADIUS Accounting agent should read the username. The default value is `#{session.logon.last.username}`.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2017-02-08 apm policy agent acct-radius(1)

apm policy agent acct-tacacsplus

NAME

acct-tacacsplus - Manages a TACACS+(r) Account agent.

MODULE

apm policy agent

SYNTAX

Configure the acct-tacacsplus component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create acct-tacacsplus [name]
modify acct-tacacsplus [name]
options
  app-service [[string] | none]
  server [[string] | none]
```

DISPLAY

```
list acct-tacacsplus
list acct-tacacsplus [ [ [name] | [glob] | [regex] ] ... ]
show running-config acct-tacacsplus
show running-config acct-tacacsplus [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  current-module
  non-default-properties
  app-service
  partition
```

DELETE

```
delete acct-tacacsplus [name]
```

DESCRIPTION

You can use the acct-tacacsplus component to configure a TACACS+ Account agent.

EXAMPLES

```
create acct-tacacsplus MyADQueryagent { server "companyAD" }
Creates the agent type TACACS+ Account named MyADQueryagent that uses the companyAD server.
```

```
list acct-tacacsplus all
Displays a list of TACACS+ Account agents and the server associated with each agent.
```

```
delete acct-tacacsplus MyADagent
Deletes the MyADagent TACACS+ Account agent.
```

OPTIONS

[name]
Specifies the name of an acct-tacacsplus agent. This setting is required.

partition
Displays the partition within which the component resides.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

server
Specifies the TACACS+ Account server that the system uses for queries and authentication.

SEE ALSO

tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

apm policy agent api-authentication

NAME

api-authentication - Manages API Authentication agent.

MODULE

apm policy agent

SYNTAX

Configure the api-authentication component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create api-authentication [name]

modify api-authentication [name]

options:

app-service [[string] | none]

edit api-authentication [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list api-authentication

list api-authentication [[[name] | [glob] | [regex]] ...]

show running-config api-authentication

show running-config api-authentication [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show api-authentication

show api-authentication [name]

DELETE

delete api-authentication [name]

DESCRIPTION

You can use the api-authentication component to create an API Authentication agent to inspect authentication header in the request and set appropriate sub-session variable.

EXAMPLES

create api-authentication my_api_authentication_agent

Creates the my_api_authentication_agent API Authentication agent.

list api-authentication

Displays a list of API Authentication agents.

delete api-authentication my_api_authentication_agent

Deletes the API Authentication agent named my_api_authentication_agent.

OPTIONS

[name]

Specifies the name of a API Authentication agent. This option is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

api-protection

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

apm policy agent api-server-selection

NAME

api-server-selection - Manages API Server Selection agent.

MODULE

apm policy agent

SYNTAX

Configure the api-server-selection component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create api-server-selection [name]

modify api-server-selection [name]

options:

app-service [[string] | none]

server [[string] | none]

edit api-server-selection [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list api-server-selection

list api-server-selection [[[name] | [glob] | [regex]] ...]

show running-config api-server-selection

show running-config api-server-selection [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show api-server-selection

show api-server-selection [name]

DELETE

delete api-server-selection [name]

DESCRIPTION

You can use the api-server-selection component to create an API Server Selection agent to assign a server config object to send the request.

EXAMPLES

create api-server-selection my_api_server_sel_ag server my_server

Creates the my_api_server_sel_ag API Server Selection agent.

list api-server-selection

Displays a list of API Server Selection agents.

delete api-server-selection my_api_server_sel_agent

Deletes the API Server Selection agent named my_api_server_sel_agent.

OPTIONS

[name]

Specifies the name of a API Server Selection agent. This option is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

server

Specifies the server config object name.

SEE ALSO

api-protection server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

apm policy agent decision-box

NAME

decision-box - Manages a Decision Box agent.

MODULE

apm policy agent

SYNTAX

Configure the decision-box component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create decision-box [name]
modify decision-box [name]
options
  app-service [[string] | none]
  customization-group [name]
```

DISPLAY

```
list decision-box
list decision-box [ [ [name] | [glob] | [regex] ] ... ]
show running-config decision-box
show running-config decision-box [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  app-service
  current-module
  non-default-properties
  one-line
  partition
```

DELETE

```
delete decision-box ([name] | all)
```

DESCRIPTION

You can use the decision-box component to configure a Decision Box agent.

EXAMPLES

```
create dynamic-acl MyADQueryagent
Creates the Decision Box agent named MyADQueryagent.
```

```
list decision-box all
Displays a list of Decision Box agents.
```

```
delete decision-box MyADagent
Deletes the MyADagent Decision Box agent.
```

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

customization-group
Specifies the name of the existing customization group to which the agent belongs.

[name]
Specifies the name of a Decision Box agent. This setting is required.

partition
Displays the partition within which the component resides.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

apm policy agent dynamic-acl

NAME

dynamic-acl - Manages a Dynamic ACL agent.

MODULE

apm policy agent

SYNTAX

Configure the dynamic-acl component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create dynamic-acl [name]
modify dynamic-acl [name]
options
  app-service [[string] | none]
  entries [ add | delete | modify | none | replace-all-with]
```

DISPLAY

```
list dynamic-acl
list dynamic-acl [ [ [name] | [glob] | [regex] ] ... ]
show running-config dynamic-acl
show running-config dynamic-acl [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  app-service
  current-module
  non-default-properties
  one-line
  partition
```

DELETE

```
delete dynamic-acl [name]
```

DESCRIPTION

You can use the dynamic-acl component to create and manage a Dynamic access control list (acl) agent that parses ACL text input with a specified format from a specified session variable, assigns the parsed entry into a Dynamic ACL object, and assigns it into a current user session. An ACL is a set of restrictions associated with a resource or favorite that defines access for users and groups.

EXAMPLES

```
create dynamic-acl { { acl [ format [f5
| cisco] ] source } } }>
Creates the Dynamic ACL agent named MyDynamicAclAgent.
```

```
list dynamic-acl
Displays a list of Dynamic ACL agents.
```

```
delete dynamic-acl MyDynamicAclAgent
Deletes the Dynamic ACL agent named MyDynamicAclAgent.
```

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

entries
Specifies the name of the entry to assign this dynamic access control list.

[name]
Specifies the name of the Dynamic Acl agent. This setting is required.

partition
Displays the partition within which the component resides.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy agent dynamic-acl(1)

apm policy agent ending-allow

NAME

ending-allow - Manages an Ending Allow agent.

MODULE

apm policy agent

SYNTAX

Configure the ending-allow component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create ending-allow [name]
modify ending-allow [name]
options:
  app-service [[string] | none]
```

DISPLAY

```
list ending-allow
list ending-allow [ [ [name] | [glob] | [regex] ] ... ]
show running-config ending-allow
show running-config ending-allow [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  app-service
  current-module
  non-default-properties
  one-line
  partition
```

DELETE

```
delete ending-allow ([name] | all)
```

DESCRIPTION

Access policy endings indicate the final outcome of a branch of an access policy. An Allow ending is a successful ending in which the system displays the user's home page and grants access to a webtop connection.

EXAMPLES

```
create ending-allow MyEndingAllowAgent { }
Creates the Ending Allow agent named MyEndingAllowAgent.
```

```
list ending-allow
Displays a list of Ending Allow agents.
```

```
delete ending-allow MyEndingAllowAgent
Deletes the Ending Allow agent named MyEndingDeniedAgent.
```

OPTIONS

[name]
Specifies the name of an Ending Allow agent. This option is required.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

partition
Displays the partition within which the component resides.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy agent ending-allow(1)

apm policy agent ending-deny

NAME

ending-deny - Manages an Ending Deny agent.

MODULE

apm policy agent

SYNTAX

Configure the ending-deny component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create ending-deny [name]
modify ending-deny [name]
options
  app-service [[string] | none]
  customization-group [name]
```

DISPLAY

```
list ending-deny
list ending-deny [ [ [name] | [glob] | [regex] ] ... ]
show running-config ending-deny
show running-config ending-deny [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  app-service
  current-module
  non-default-properties
  one-line
  partition
```

DELETE

```
delete ending-deny ([name] | all)
```

DESCRIPTION

Access policy endings indicate the final outcome of a branch of an access policy. The Logon Deny ending is the final result of an unsuccessful logon attempt (the failure could be caused by an incorrect logon attempt, a security requirement incompatibility, or the use of an unsupported device). Upon reaching a Logon Deny ending, the user sees an error message. You can use the ending-deny component to create and manage an Ending Deny agent.

EXAMPLES

```
create ending-deny MyEndingDenyAgent customization-group MyLogOffCG
Creates the Ending Deny agent named MyEndingDenyAgent that is associated with the MyLogOffCG customization group.
```

```
list ending-deny
Displays a list of Ending Deny agents.
```

```
delete ending-deny MyEndingDenyAgent
Deletes the Ending Deny agent named MyEndingDenyAgent.
```

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

customization-group
Specifies the name of the existing customization-group to which the agent belongs. It enables you to customize the logon deny page. For example, you can indicate a specific reason for the denial of access. This setting is required, and the customization group that you assign must be of the type logout.

[name]
Specifies the name of an Ending Deny agent. This setting is required.

partition
Displays the partition within which the component resides.

SEE ALSO

tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy agent ending-deny(1)

NAME
ending-redirect - Manages an Ending Redirect agent.

MODULE
apm policy agent

SYNTAX
Configure the ending-redirect component within the policy agent module using the following syntax.

CREATE/MODIFY
create ending-redirect [name]
modify ending-redirect [name]
options
app-service [[string] | none]
close-session [true | false]
url [value]

DISPLAY
list ending-redirect
list ending-redirect [[[name] | [glob] | [regex]] ...]
show running-config ending-redirect
show running-config ending-redirect [[[name] | [glob] | [regex]] ...]
options:
all
all-properties
app-service
current-module
non-default-properties
one-line
partition

DELETE
delete ending-redirect ([name] | all)

DESCRIPTION
The Redirect ending can be used to redirect the user, rather than allowing or denying a connection. It can also send a user directly to an update script or to different server or landing URI. Upon reaching a Redirect ending, the user sees a screen indicating that they are being redirected to a different URL. You can use the ending-redirect component to create and manage an Ending Redirect agent.

EXAMPLES
create ending-redirect MyEndingRedirectAgent { url "http://www.myweb.com" }
Creates the Ending Redirect agent named MyEndingRedirectAgent that redirects a connection to http://www.myweb.com.

=item

Creates an agent using the current protocol and the session variable %`{session.server.network.protocol}`

list ending-redirect
Displays a list of Ending Redirect agents.

delete ending-redirect MyEndingRedirectAgent
Deletes the Ending Redirect agent named MyEndingRedirectAgent.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

close-session
Redirects to the specified URI after closing the session if enabled. Otherwise, redirect to the specified URI without closing the session. The default is enabled.

[name]
Specifies the name of an Ending Redirect agent. This option is required.

url Specifies the URL to which the system redirects the original request. This option is required, and you must specify an absolute URL.

An absolute URL specifies the exact location of a file or directory on the Internet.

SEE ALSO
tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm policy agent ending-redirect(1)

apm policy agent endpoint-check-machine-cert

NAME

endpoint-check-machine-cert - Manages an End-point Check Machine certificate agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-check-machine-cert component within the apm policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-check-machine-cert [name]

modify endpoint-check-machine-cert [name]

options:

allow-elevation [true| false]

app-service [[string] | none]

ca-profile-name [value]

issuer [value]

save-cert [true| false]

serial-number [integer]

store-location [machine | user]

store-name [value]

subject-alt-name [value]

subject-match-fqdn [value]

edit endpoint-check-machine-cert [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list endpoint-check-machine-cert

list endpoint-check-machine-cert [[[name] | [glob] | [regex]] ...]

show running-config endpoint-check-machine-cert

show running-config endpoint-check-machine-cert [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

app-service

current-module

non-default-properties

one-line

partition

DELETE

delete endpoint-check-machine-cert [name]

DESCRIPTION

Endpoint security is a centrally-managed method of monitoring and maintaining client-system security.

The endpoint-check-machine-cert component checks for the presence of a valid machine certificate on Windows/Mac client systems during access policy validation.

EXAMPLES

create endpoint-check-machine-cert MyMCagent

Creates the Endpoint Check Machine certificate agent named MyMCagent in the Common partition.

list endpoint-check-machine-cert

Displays a list of Endpoint Check Machine certificate agents.

delete endpoint-check-machine-cert MyMCagent

Deletes the MyMCagent Endpoint Check Machine certificate agent.

OPTIONS

allow-elevation

Specifies that User Account Control (UAC) prompts are allowed during private key checking.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

ca-profile-name

Specifies the name of the certificate authority profile to validate the certificate.

issuer

Specifies the name used to match the issuer name in the machine certificate.

[name]

Specifies the name of an external logon page agent. This option is required.

partition

Displays the partition within which the component resides.

save-cert

Specifies to store the entire machine certificate in a session variable.

serial-number

Specifies the serial number used to match the serial number of the machine certificate.

store-location

Specifies the location of the certificate store on the client machine.

store-name

Specifies the name of the certificate store on the client machine.

subject-alt-name

Specifies the name used to match the subject-alt-name in the machine certificate.

subject-match-fqdn

Specifies if lookup must match fully qualified domain name (FQDN) in the machine certificate.

SEE ALSO

apm policy agent endpoint-check-software, apm policy agent endpoint-linux-check-file, apm policy agent endpoint-linux-check-process, apm policy agent endpoint-mac-check-file, apm policy agent endpoint-mac-check-process, apm policy agent endpoint-windows-check-file, apm policy agent endpoint-windows-browser-cache-cleaner, apm policy agent endpoint-windows-check-process, apm policy agent endpoint-windows-check-registry, apm policy agent endpoint-windows-info-os, apm policy agent endpoint-machine-info, apm policy agent endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm policy agent endpoint-check-machine-cert(1)

apm policy agent endpoint-check-software

NAME

endpoint-check-software - Manages an Endpoint Software Check agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-check-software component within the apm policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-check-software [name]

modify endpoint-check-software [name]

options:

collect [true | false]

continuous-check [true | false]

type [antivirus | firewall | patch-management | antispyware | peer-to-peer | hard-disk-encryption | health-agent]

check-list-type [required | allow | deny]

items [vendor_id | product_id | state | version | db-age | db-version | last-scan | missing-updates | platform]

edit endpoint-check-software [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list endpoint-check-software

list endpoint-check-software [[[name] | [glob] | [regex]] ...]

show running-config endpoint-check-software

show running-config endpoint-check-software [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

app-service

current-module

non-default-properties

one-line

partition

DELETE
delete endpoint-check-software ([name] | all)

DESCRIPTION

Endpoint security is a centrally-managed method of monitoring and maintaining client-system security. You can use the endpoint-check-software component to create and manage an agent that enforces monitoring of various client-system security third party software. Different types of third party software supported are described below in options.

The configuration attributes in the items option are generic and therefore for a given software type only certain items attributes are useful, rest of the attributes are ignored even if they are configured. For example: for type=peer-to-peer only vendor_id, product_id, state and version are considered and rest of the items like db-age, db-version etc are ignored. Following is the list of useful attributes corresponding to the software type:

Common to all software type:
vendor_id, product_id, version, platform, state

antivirus & antispware:
db-age, db-version, last-scan

patch-management:
missing-updates

EXAMPLES

create endpoint-check-software MyEndpointWCagent items state enabled add
Creates the Endpoint Check Software agent named MyEndpointWCagent, which verifies that the specified third party software on the client is compliant with system administrators configuration, which may just check for the installation or monitor the state of the software

list endpoint-check-software
Displays a list of Endpoint Software Check agents.

delete endpoint-check-software MyEndpointWCagent
Deletes the Endpoint Software Check agent named MyEndpointWCagent.

OPTIONS

items
Adds items to or deletes items from an Endpoint Software Check agent. You can specify the following attributes for the software:

check-list-type Specifies how the list of software should be checked
required:

Client is required to have at least one of the software configured in the list in order to pass the access policy. And that software should satisfy all the configuration fields e.g. state, version etc.

allow: Client is allowed to have any of the software configured in the list but NOT any other than that, in order to pass the access policy. List is treated as whitelist. A given client software will not match unless it satisfies all the configuration fields (e.g. state, version etc). NOTE: The check will also be successful if client has no software installed at all. List of software is treated as whitelist.

deny: Client should NOT have any software configured in the list in order to pass the access policy. And that software should satisfy all the configuration fields (e.g. state, version etc). NOTE: The check will also be successful if client has no software installed at all. List of software is treated as blacklist.

db-age
Specifies the maximum age of the anti-virus/anti-spyware database that you want an Endpoint Software Check agent to verify the presence of on the client in order to allow the access policy to pass.

db-version
Specifies the version of the anti-virus/anti-spyware database that you want an Endpoint Software Check agent to verify the presence of on the client in order to allow the access policy to pass.

product_id
Specifies the product ID of the software that you want an Endpoint Software Check agent to verify the presence of on the client in order to allow the access policy to pass.

vendor_id
Specifies the vendor ID of the software that you want an Endpoint Software Check agent to verify the presence of on the client in order to allow the access policy to pass.

NOTE: If none of the vendor id or product id is defined then check is performed for any of the software of given type. If both vendor id and product id are configured then, product id is ignored and only vendor id is considered. Vendor ID always takes precedence. A vendor can have many products. Each product (of every vendor) has unique ID assigned to them. Similarly, every vendor is assigned a unique ID too which is separate from product ID. If you want to check every software from a vendor then specify vendor_id only.

state
State means different things to different software type. The state can be enabled, disabled or unspecified. The default is unspecified.

antivirus and antispware:
When the state is set to enabled or disabled, agent verifies that the specified

antivirus/antispysware software has real time protection enabled or disabled on the client that is attempting to connect. When state is unspecified, it ignores the state.

patch-management:

When the state is set to enabled, agent verifies that the specified PM software is running on the client that is attempting to connect. When its set to unspecified, state of the software is ignored.

firewall:

When the state is enabled or disabled, agent verifies that the specified firewall software has real time protection enabled or disabled on the client that is attempting to connect. When state is unspecified, the software state is ignored.

peer-to-peer:

When the state is set to enabled agent verifies that the peer-to-peer software is running on the client that is attempting to connect. When state is unspecified, the agent only verifies that the software is installed or not.

hard-disk-encryption:

When the state is set to enabled agent verifies that all disk volumes are encrypted on the client that is attempting to connect. When the state is set to disabled agent verifies that system disk volume is encrypted on the client that is attempting to connect. When state is unspecified, the agent only verifies that the software is installed or not.

health-agent:

When the state is set to enabled agent verifies that endpoint client is compliant with the health policy set out by the site administrator.

version

Specifies the version of the software that you want an Endpoint Software Check agent to verify the presence of on the client in order to allow the access policy to pass.

last-scan

Specifies the maximum allowed duration without the full system scan of endpoint client that software agent can accept in order to allow the access policy to pass. It is specified in number of days.

missingupdates

Specifies the maximum number of allowed missing critical updates of the PM software at the endpoint client in order to allow the access policy to pass. Leave blank to ignore number of missing critical updates. Specify 0 to make sure endpoint client is up-to-date

platform

Specifies the platform. It could be any of the following: windows, linux, mac or any. The default is any.

type Its the type of the third party software to be monitored on the client system. It could be any of the following: antivirus, firewall, patch-management, antispysware, peer-to-peer, hard-disk-encryption, health-agent

collect

This setting is ignored.

continuous-check

Continuously check the items, and end the session if the result changes. The default is false.

[name]

Specifies the name of an Endpoint Software Check agent. This option is required.

partition

Displays the partition within which the component resides.

SEE ALSO

apm policy agent endpoint-linux-check-file, apm policy agent endpoint-linux-check-process, apm policy agent endpoint-mac-check-file, apm policy agent endpoint-mac-check-process, apm policy agent endpoint-windows-browser-cache-cleaner, apm policy agent endpoint-windows-check-file, apm policy agent endpoint-check-machine-cert, apm policy agent endpoint-windows-check-process, apm policy agent endpoint-windows-check-registry, apm policy agent endpoint-windows-group-policy, apm policy agent endpoint-windows-info-os, apm policy agent endpoint-machine-info, apm policy agent endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 apm policy agent endpoint-check-software(1)

NAME

endpoint-linux-check-file - Manages an Endpoint Linux Check File agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-linux-check-file component within the policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-linux-check-file [name]

modify endpoint-linux-check-file [name]

options:

continuous-check [true | false]

app-service [[string] | none]

files [filename | md5 | modified | size]

edit endpoint-linux-check-file [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list endpoint-linux-check-file

list endpoint-linux-check-file [[[name] | [glob] | [regex]] ...]

show running-config endpoint-linux-check-file

show running-config endpoint-linux-check-file [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

app-service

current-module

non-default-properties

one-line

partition

DELETE

delete endpoint-linux-check-file ([name] | all)

DESCRIPTION

Access Policy Manager checks for the presence of one or more files on a client that is attempting to connect.

If a file with the described properties exists, the action goes to the successful branch. If the file does not exist, or a file exists but one or more properties are not correct, the action goes to the fallback branch.

You can use the endpoint-linux-check-file component to create or manage an Endpoint Linux Check File agent that verifies the presence of specified Linux files on a client.

EXAMPLES

```
create endpoint-linux-check-file Myprofile_act_file_check_ag { files { filename "/tmp/demo/demofile" md5
"6b61ad518c23650b17e738e1fa2bb04e" modified 2007-06-01 10:30:10 size 12 } { filename "/tmp/demo/testfile" md5
"f20d9f2072bbeb6691c0f9c5099b01f3" size 9 } }
```

Creates the Endpoint Linux Check File agent named Myprofile_act_file_check_ag that checks that the client contains two files located in the /tmp/demo directory: a 12 byte file named demofile that was modified no later than January 6, 2007 at 10:30 and has an MD5 checksum of 6b61ad518c23650b17e738e1fa2bb04e, and a 9-byte file named testfile that has an MD5 check sum of f20d9f2072bbeb6691c0f9c5099b01f3.

```
list all endpoint-linux-check-file Company8profile_act_file_check_ag
```

Displays information about the Endpoint Linux Check File agent named Company8profile_act_file_check_ag.

```
delete endpoint-linux-check-file Company8profile_act_check_file { files { filename "/tmp/demo/demofile" } }
```

Deletes the /tmp/demo/demofile file from the Endpoint Linux Check File agent named Company8profile_act_file_check_ag.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

files

Adds files to or deletes files from an Endpoint Linux Check File agent. You can specify the following attributes of the files that you want an Endpoint Linux Check File agent to verify the presence of on the client in order to allow the access policy to pass.

filename

Specifies the name of the file and includes the full path. The Endpoint linux Check File agent that you are creating must be able to verify the file's presence on the client for the access policy to pass. When you add a file to or delete a file from the agent, this setting is required.

md5 Specifies the value of an MD5 checksum. The Endpoint Linux Check File agent you are creating must be able to match the checksum on the client for the access policy to pass. The default is none.

modified

Specifies the last modified date of the specified file. The Endpoint Linux Check File agent you are creating must verify this date on the client for the access policy to pass. The default is 1970-01-01 00:00:00.

size Specifies the size, in bytes, of the specified file. The Endpoint Linux Check File agent you are creating must verify this size on the client for the access policy to pass. The default is 0.

continuous-check
Continuously check the files, and end the session if the result changes. The default is false.

[name]
Specifies the name of an Endpoint Linux Check File agent. This setting is required.

partition
Displays the partition within which the component resides.

SEE ALSO

endpoint-check-software, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 apm policy agent endpoint-linux-check-file(1)

apm policy agent endpoint-linux-check-process

NAME
endpoint-linux-check-process - Manages an Endpoint Linux Check Process agent.

MODULE
apm policy agent

SYNTAX
Configure the endpoint-linux-check-process component within the policy agent module using the following syntax.

CREATE/MODIFY
create endpoint-linux-check-process [name]
modify endpoint-linux-check-process [name]
options
continuous-check [true | false]
app-service [[string] | none]
expression [string | none]

edit endpoint-linux-check-process [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list endpoint-linux-check-process
list endpoint-linux-check-process [[[name] | [glob] | [regex]] ...]
show running-config endpoint-linux-check-process
show running-config endpoint-linux-check-process [[[name] | [glob] | [regex]] ...]
options:
all
all-properties
app-service
current-module
non-default-properties
one-line
partition

DELETE
delete endpoint-linux-check-process [name]

DESCRIPTION
You can use the endpoint-linux-check-process component to create and manage an Endpoint Linux Check Process agent that collects information about the Linux processes running on the client.

EXAMPLES
create endpoint-linux-check-process MyEndpointWCPagent { (bash OR top) AND firefox }
Creates the Endpoint Linux Check Process agent named MyEndpointWCPagent that checks that the client has either bash or top, and firefox launched.

list endpoint-linux-check-process

Displays a list of Endpoint Linux Check Process agents.

delete endpoint-linux-check-process MyEndpointWCPagent
Deletes the Endpoint Linux Check Process agent named MyEndpointWCPagent.

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

expression
Specifies the expression that you want an Endpoint Linux Check Process agent to use to verify the processes that are running on the client to allow the access policy to pass. You can use the following operators: AND, OR, NOT, (and). You can use wildcards in the process name, for example, navapvc.*.

If the check is successful, the system returns 1. If the check fails, the system returns 0. If the expression is incorrect, the system returns -1.

continuous-check
Continuously check the expression, and end the session if the result changes. The default is false.

[name]
Specifies the name of an Endpoint Linux Check Process agent. This setting is required.

partition
Displays the partition within which the component resides.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-group-policy, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-05-30 apm policy agent endpoint-linux-check-process(1)

apm policy agent endpoint-mac-check-file

NAME
endpoint-mac-check-file - Manages an Endpoint Macintosh Check File agent.

MODULE
apm policy agent

SYNTAX
Configure the endpoint-mac-check-file component within the policy agent module using the following syntax.

CREATE/MODIFY
create endpoint-mac-check-file [name]
modify endpoint-mac-check-file [name]
options
continuous-check [true | false]
app-service [[string] | none]
files [filename | md5 | modified | size]

edit endpoint-mac-check-file [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list endpoint-mac-check-file
list endpoint-mac-check-file [[[name] | [glob] | [regex]] ...]
show running-config endpoint-mac-check-file
show running-config endpoint-mac-check-file [[[name] | [glob] | [regex]] ...]
options:
all
all-properties
app-service
current-module
non-default-properties
one-line

partition

DELETE

delete endpoint-mac-check-file ([name] | all)

DESCRIPTION

Access Policy Manager checks for the presence of one or more files on a client that is attempting to connect. If a file with the described properties exists, the action goes to the successful branch. If the file does not exist, or a file exists but one or more properties are not correct, the action goes to the fallback branch.

You can use the endpoint-mac-check-file component to create or manage an Endpoint Macintosh Check File agent that verifies the presence of specified Macintosh files on a client.

EXAMPLES

```
create endpoint-mac-check-file Myprofile_act_file_check_ag { files { filename "/tmp/demo/demofile" md5
"6b61ad518c23650b17e738e1fa2bb04e" modified 2007-06-01 10:30:10 size 12 } { filename "/tmp/demo/testfile" md5
"f20d9f2072bbeb6691c0f9c5099b01f3" size 9 } }
```

Creates the Endpoint Macintosh Check File agent named Myprofile_act_file_check_ag that checks that the client contains two files located in the /tmp/demo directory: a 12 byte file named demofile that was modified no later than January 6, 2007 at 10:30 and has an MD5 checksum of 6b61ad518c23650b17e738e1fa2bb04e, and a 9 byte file named testfile that has an MD5 check sum of f20d9f2072bbeb6691c0f9c5099b01f3.

```
list all endpoint-mac-check-file Company8profile_act_file_check_ag
```

Displays information about the Endpoint Macintosh Check File agent named Company8profile_act_file_check_ag.

```
delete endpoint-mac-check-file Company8profile_act_check_file { files { filename "/tmp/demo/demofile" } }
```

Deletes the /tmp/demo/demofile file from the Endpoint Macintosh Check File agent named Company8profile_act_file_check_ag.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

files

Adds files to or deletes files from an Endpoint Macintosh Check File agent. You can specify the following attributes of the files that you want an Endpoint Macintosh Check File agent to verify the presence of on the client to allow the access policy to pass:

filename

Specifies the name of the file and includes the full path. The Endpoint Macintosh Check File agent that you are creating must be able to verify the file's presence on the client for the access policy to pass. When you add a file to or delete a file from the agent, this setting is required.

md5 Specifies the value of an MD5 checksum. The Endpoint Macintosh Check File agent you are creating must be able to match the checksum on the client for the access policy to pass. The default is none.

modified

Specifies the last modified date of the specified file. The Endpoint Macintosh Check File agent you are creating must verify this date on the client for the access policy to pass. The default is 1970-01-01 00:00:00.

size Specifies the size, in bytes, of the specified file. The Endpoint Macintosh Check File agent you are creating must verify this size on the client for the access policy to pass. The default is 0.

continuous-check

Continuously check the files, and end the session if the result changes. The default is false.

[name]

Specifies the name of an Endpoint Macintosh Check File agent. This setting is required.

partition

Displays the partition within which the component resides.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-group-policy, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.AAA

BIG-IP 2013-05-30 apm policy agent endpoint-mac-check-file(1)

apm policy agent endpoint-mac-check-process

NAME

endpoint-mac-check-process - Manages an Endpoint Macintosh Check Process agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-mac-check-process component within the policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-mac-check-process [name]

modify endpoint-mac-check-process [name]

options

continuous-check [true | false]

app-service [[string] | none]

expression [string | none]

edit endpoint-mac-check-process [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list endpoint-mac-check-process

list endpoint-mac-check-process [[[name] | [glob] | [regex]] ...]

show running-config endpoint-mac-check-process

show running-config endpoint-mac-check-process [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

app-service

current-module

non-default-properties

one-line

partition

DELETE

delete endpoint-mac-check-process ([name] | all)

DESCRIPTION

You can use the endpoint-mac-check-process component to create and manage an Endpoint Macintosh Check Process agent that collects information about the Macintosh processes running on the client.

EXAMPLES

```
create endpoint-mac-check-process MyEndpointWCPagent { (bash OR top) AND firefox }
```

Creates the Endpoint Macintosh Check Process agent named MyEndpointWCPagent that checks that the client has either bash or top, and firefox launched.

```
list endpoint-mac-check-process
```

Displays a list of Endpoint Macintosh Check Process agents.

```
delete endpoint-mac-check-process MyEndpointWCPagent
```

Deletes the Endpoint Macintosh Check Process agent named MyEndpointWCPagent.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

expression

Specifies the expression that you want an Endpoint Macintosh Check Process agent to use to verify the processes that are running on the client in order to allow the access policy to pass. You can use the following operators: AND, OR, NOT, (and). You can use wildcards in the process name, for example, navapvc.*.

If the check is successful, the system returns 1. If the check fails, the system returns 0. If the expression is incorrect, the system returns -1.

continuous-check

Continuously check the expression, and end the session if the result changes. The default is false.

[name]

Specifies the name of an Endpoint Macintosh Check Process agent. This setting is required.

partition

Displays the partition within which the component resides.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-29 apm policy agent endpoint-mac-check-process(1)

apm policy agent endpoint-machine-info

NAME

endpoint-machine-info - Manages an Endpoint Machine Information agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-machine-info component within the policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-machine-info [name]

modify endpoint-machine-info [name]

options:

app-service [[string] | none]

edit endpoint-machine-info [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DISPLAY

list endpoint-machine-info

list endpoint-machine-info [[[name] | [glob] | [regex]] ...]

show running-config endpoint-machine-info

show running-config endpoint-machine-info [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show endpoint-machine-info

show endpoint-machine-info [name]

DELETE

delete endpoint-machine-info [name]

DESCRIPTION

You can use the endpoint-machine-info component to create and manage an agent that collects information about the machine that is attempting to connect.

OPTIONS

[name]

Specifies the name of the an Endpoint Check Machine Information agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-group-policy, endpoint-windows-info-os, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-05-30 apm policy agent endpoint-machine-info(1)

apm policy agent endpoint-windows-browser-cache-cleaner

NAME

endpoint-windows-browser-cache-cleaner - Manages an Endpoint Windows Browser Cache Cleaner agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-windows-browser-cache-cleaner component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create endpoint-windows-browser-cache-cleaner [name]
```

```
modify endpoint-windows-browser-cache-cleaner [name]
```

options

app-service [[string] | none]

cache-clean-type [all | all-except-css-js | all-except-img-css-js | none]

clean-passwords [false | true]

empty-recycle-bin [false | true]

idle-timeout [| immediate | indefinite]

idle-timeout-screen-lock []

monitor-webtop [enable | disable]

partition

remove-connection-entry [false | true]

```
edit endpoint-windows-browser-cache-cleaner [ [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties

non-default-properties

DISPLAY

```
list endpoint-windows-browser-cache-cleaner
```

```
list endpoint-windows-browser-cache-cleaner [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config endpoint-windows-browser-cache-cleaner
```

```
show running-config endpoint-windows-browser-cache-cleaner [ [ [name] | [glob] | [regex] ] ... ]
```

options:

all

all-properties

app-service

current-module

non-default-properties

one-line

partition

DELETE

```
delete endpoint-windows-browser-cache-cleaner ([name] | all)
```

DESCRIPTION

Endpoint security is a centrally-managed method of monitoring and maintaining client-system security. You can use the endpoint-windows-browser-cache-cleaner component to create and manage an Endpoint Windows Browser Cache Cleaner agent. This agent cleans items from the client browser and computer after logoff, and also enforces session inactivity timeouts.

EXAMPLES

```
create endpoint-windows-browser-cache-cleaner MyEndpointWBCCagent idle timeout 0
```

Creates the Endpoint Windows Browser Cache Cleaner agent named MyEndpointWBCCagent that does not enforce a timeout.

```
create endpoint-windows-browser-cache-cleaner MyEndpointWBCCagent { idle timeout 0 clean passwords enable }
```

Creates the Endpoint Windows Browser Cache Cleaner agent named MyEndpointWBCCagent that does not enforce a timeout, but does clear saved passwords from the client after logoff.

```
list endpoint-windows-browser-cache-cleaner
```

Displays a list of Endpoint Windows Browser Cache Cleaner agents.

```
delete endpoint-windows-browser-cache-cleaner MyEndpointWBCCagent
```

Deletes the Endpoint Windows Browser Cache Cleaner agent named MyEndpointWBCCagent.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

cache-clean-type

Specifies which browser cache temporary files are removed. If set to all, the temporary files are removed. If set to all-except-css-js, the browser cache is cleared, but all style sheets and JavaScript are left on the browser cache. If set to all-except-img-css-js, the browser cache is cleared, but all style sheets, JavaScript, and images are left on the browser cache. The default is all.

clean-passwords

When true, the Endpoint Windows Browser Cache Cleaner agent ensures that saved passwords are cleared from the client after logoff. The default is false.

empty-recycle-bin

When true, the Endpoint Windows Browser Cache Cleaner agent empties the Recycle Bin on the client after logoff. The default is false.

idle-timeout

Specifies the number of minutes that the client session can be idle before the Endpoint Windows Browser Cache Cleaner agent disconnects the session. The default is 0, which enforces no timeout. This is a required setting.

idle-timeout-screen-lock

Specifies the number of minutes the system can receive no user input before the workstation is locked. The default is 0, which specifies no timeout enforced.

monitor-webtop

When true, the Endpoint Windows Browser Cache Cleaner agent forces session termination if the browser or webtop is closed. The default is false.

[name]

Specifies the name of the Endpoint Windows Browser Cache Cleaner agent. This is a required setting.

partition

Displays the partition within which the component resides.

remove-connection-entry

When true, the Endpoint Windows Browser Cache Cleaner agent removes the connection from the Network Connections Dial-up Networking folder on the client. The default is false.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2016-apm1policy agent endpoint-windows-browser-cache-cleaner(1)

apm policy agent endpoint-windows-check-file

NAME

endpoint-windows-check-file - Manages an Endpoint Windows Check File agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-windows-check-file component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create endpoint-windows-check-file [name]
modify endpoint-windows-check-file [name]
options
  continuous-check [ true | false ]
  app-service [[string] | none]
  files [ filename | md5 | modified | operation | signer | size | version ]
```

```
edit endpoint-windows-check-file [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list endpoint-windows-check-file
list endpoint-windows-check-file [ [ [name] | [glob] | [regex] ] ... ]
show running-config endpoint-windows-check-file
show running-config endpoint-windows-check-file [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  app-service
  current-module
```

non-default-properties
one-line
partition

DELETE

delete endpoint-windows-check-file ([name] | all)

DESCRIPTION

Access Policy Manager checks for the presence of one or more files on a client that is attempting to connect. If a file with the described properties exists, the action goes to the successful branch. If the file does not exist, or a file exists but one or more properties are not correct, the action goes to the fallback branch.

You can use the endpoint-windows-check-file component to create or manage an Endpoint Windows Check File agent that verifies the presence of specified Windows files on a client.

EXAMPLES

```
create endpoint-windows-check-file Myprofile_act_file_check_ag { files { filename "C:\\demo\\demofile" md5
"6b61ad518c23650b17e738e1fa2bb04e" modified 2007-06-01 10:30:10 size 12 } { filename "C:\\demo\\test.file" md5
"f20d9f2072bbeb6691c0f9c5099b01f3" size 9 } }
```

Creates the Endpoint Windows Check File agent named Myprofile_act_file_check_ag that checks that the client contains two files located in the C:\demo directory: a 12 byte file named demofile that was modified no later than January 6, 2007 at 10:30 and has an MD5 checksum of 6b61ad518c23650b17e738e1fa2bb04e, and a 9 byte file named test.file that has an MD5 check sum of f20d9f2072bbeb6691c0f9c5099b01f3.

```
list all endpoint-windows-check-file Company8profile_act_file_check_ag
```

Displays information about the Endpoint Windows Check File agent named Company8profile_act_file_check_ag.

```
delete endpoint-windows-check-file Company8profile_act_check_file { files { filename "C:\\demo\\demofile" } }
```

Deletes the C:\demo\demofile file from the Endpoint Windows Check File agent named Company8profile_act_file_check_ag.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

files

Adds files to or deletes files from an Endpoint Windows Check File agent. You can specify the following attributes for the files that you want an Endpoint Windows Check File agent to verify the presence of on the client to allow the access policy to pass.

filename

Specifies a file name and includes the full path. The Endpoint windows Check File agent you are creating must be able to verify the file's presence on the client for the access policy to pass.

When you add a file to or delete a file from the agent, this setting is required.

md5 Specifies the value of an MD5 checksum. The Endpoint windows Check File agent that you are creating must match the checksum on the client for the access policy to pass. The default is none.

modified

Specifies the last modified date of the specified file. The Endpoint windows Check File agent you are creating must verify this date on the client for the access policy to pass. The default is 1970-01-01 00:00:00.

operation

Specifies the operator that you want your Endpoint Windows Check File agent to use when verifying the attributes of the specified file on the client. The default is equal.

signer

Specifies that the Endpoint Windows Check File agent must verify that the specified file on the client is signed for the access policy to pass. The default is none.

size Specifies the size, in bytes, of the specified file. The Endpoint Windows Check File agent you are creating must verify this file size on the client for the access policy to pass. The default is 0.

version

Specifies the version of the specified file that you want your Endpoint Windows Check File agent to verify on the client for the access policy to pass. Specify the version using the following form: x.x.x.x. The maximum value is 65535.65535.65535.65535. The default is none.

continuous-check

Continuously check the files, and end the session if the result changes. The default is false.

[name]

Specifies the name of an Endpoint Windows Check File agent. This option is required.

partition

Displays the partition within which the component resides.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-check-file, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-group-policy, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 apm policy agent endpoint-windows-check-file(1)

apm policy agent endpoint-windows-check-process

NAME

endpoint-windows-check-process - Manages an Endpoint Windows Check Process agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-windows-check-process component within the policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-windows-check-process [name]

modify endpoint-windows-check-process [name]

options:

continuous-check [true | false]

app-service [[string] | none]

expression (| none)

edit endpoint-windows-check-process [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list endpoint-windows-check-process

list endpoint-windows-check-process [[[name] | [glob] | [regex]] ...]

show running-config endpoint-windows-check-process

show running-config endpoint-windows-check-process [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show endpoint-windows-check-process

show endpoint-windows-check-process [name]

DELETE

delete endpoint-windows-check-process [name]

DESCRIPTION

You can use the endpoint-windows-check-process component to create and manage an agent that collects information about the Windows processes running on the client.

EXAMPLES

```
create endpoint-windows-check-process MyEndpointWCPagent { (NISUM.exe OR blackd.exe) AND navapvc.* }
```

Creates the Endpoint Windows Check Process agent named MyEndpointWCPagent that checks that the client has either NISUM.exe or blackd.exe, and navapvc.* installed.

```
list endpoint-windows-check-process
```

Displays a list of Endpoint Windows Check Process agents.

```
delete endpoint-windows-check-process MyEndpointWCPagent delete
```

Deletes the Endpoint Windows Check Process agent named MyEndpointWCPagent.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

expression

Specifies the expression that you want an Endpoint Windows Check Process agent to use to verify the processes that are running on the client in order to allow the access policy to pass. You can use the following operators: AND, OR, NOT, (and). You can use wildcards in the process name, for example, navapvc.*.

If the check is successful, the system returns 1. If the check fails, the system returns 0. If the expression is incorrect, the system returns -1.

continuous-check

Continuously check the expression, and end the session if the result changes. The default is false.

[name]

Specifies the name of an Endpoint Windows Check Process agent. This setting is required.

partition

Displays the partition within which the component resides.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-registry, endpoint-windows-group-policy, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-05-30 apm policy agent endpoint-windows-check-process(1)

apm policy agent endpoint-windows-check-registry

NAME

endpoint-windows-check-registry - Manages an Endpoint Windows Check Registry agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-windows-check-registry component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create endpoint-windows-check-registry [name]
modify endpoint-windows-check-registry [name]
options:
  continuous-check [ true | false ]
  app-service [[string] | none]
  expression [[string] | none]
```

```
edit endpoint-windows-check-registry [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list endpoint-windows-check-registry
list endpoint-windows-check-registry [ [ [name] | [glob] | [regex] ] ... ]
show running-config endpoint-windows-check-registry
show running-config endpoint-windows-check-registry [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition
```

```
show endpoint-windows-check-registry
show endpoint-windows-check-registry [name]
```

DELETE

```
delete endpoint-windows-check-registry [name]
```

DESCRIPTION

You can use the endpoint-windows-check-registry component to create and manage an agent that collects information about the Windows registry keys on the client that is attempting to connect.

EXAMPLES

```
create endpoint-windows-check-registry MyEndpointWCRagent
{"\"HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Internet Explorer\\.\"Version\"=\"5.0.2800.0\""}
Creates the Endpoint Windows Check Registry agent named MyEndpointWCRagent that checks the registry on the client for version 5.0.2800.0 of Internet Explorer in the HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft directory.
```

```
create endpoint-windows-check-registry MyEndpointWCRagent
{"\"HKEY_LOCAL_MACHINE64\\SOFTWARE\\Microsoft\\Internet Explorer\\.\"Version\"=\"5.0.2800.0\""}

```

Creates the Endpoint Windows Check Registry agent named MyEndpointWCRagent that checks the registry on the client for version 5.0.2800.0 of Internet Explorer in the HKEY_LOCAL_MACHINE64\SOFTWARE\Microsoft directory.

Note that the registry value HKEY_LOCAL_MACHINE64 is one of the 32 and 64-bit registry keys that you can specify on 64-bit Windows versions. On 64-bit Windows systems, you can check for registry keys in either the 64-bit registry or the 32-bit registry. To specify the registry to check, append a number to the registry root key name. The following key names are supported:

HKEY_CURRENT_USER
HKEY_CURRENT_USER32
HKEY_CURRENT_USER64
HKEY_LOCAL_MACHINE
HKEY_LOCAL_MACHINE32
HKEY_LOCAL_MACHINE64
HKEY_CLASSES_ROOT
HKEY_CLASSES_ROOT32
HKEY_CLASSES_ROOT64
HKEY_USERS
HKEY_USERS32

HKEY_USERS64 HKEY values specified with a 32 allow you to check values in the 32-bit view of 64-bit registry. This is the perspective used by 32-bit applications running with on a 64-bit operating system.

HKEY values with a 64 appended allow you to check values in the 64-bit view of the registry. This is the perspective used by native 64-bit applications. When checking values on 32-bit Windows, the number of bits specified in the registry key name is ignored.

list endpoint-windows-check-registry
Displays a list of Endpoint Windows Check Registry agents.

delete endpoint-windows-check-registry MyEndpointWCRagent delete
Deletes the Endpoint Windows Check Registry agent named MyEndpointWCRagent.

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

expression
Specifies the expression that you want an Endpoint Windows Check Registry agent to use to verify the registry entries that are present on the client in order to allow the access policy to pass. You can use the following operators: AND, OR, NOT, (and).

If the check is successful, the system returns 1. If the check fails, the system returns 0. If the expression is incorrect, the system returns -1.

Important: You must use quotation marks (" ") around key and value arguments, and in data when the content contains spaces, commas, slashes, tabs, or other delimiters. If quotation marks exist as part of a registry path or value name, you must use quotation marks around those quotation marks.

Tip: The system treats data in the formats "d.d[.d][.d]" or "d,d[.d][.d]" (where d is a number) as a version number. The system treats data in the format "mm/dd/yyyy" as a date.

continuous-check
Continuously check the expression, and end the session if the result changes. The default is false.

[name]
Specifies the name of the an Endpoint Windows Check Registry agent. This option is required.

partition
Displays the partition within which the component resides.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-group-policy, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-05-30 apm policy agent endpoint-windows-check-registry(1)

NAME

endpoint-windows-group-policy - Manages an Endpoint Windows Group Policy agent.

MODULE

apm policy agent

SYNTAX

Warning: This page is obsolete. Windows Group Policy is no longer supported.

Configure the external-logon-page component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create endpoint-windows-group-policy [name]
modify endpoint-windows-group-policy [name]
options:
  app-service [[string] | none]
  policy-file { [name] }
```

```
edit endpoint-windows-group-policy [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list endpoint-windows-group-policy
list endpoint-windows-group-policy [ [ [name] | [glob] | [regex] ] ... ]
show running-config endpoint-windows-group-policy
show running-config endpoint-windows-group-policy [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition
```

```
show endpoint-windows-group-policy
show endpoint-windows-group-policy [name]
```

DELETE

```
delete endpoint-windows-group-policy [name]
```

DESCRIPTION

Endpoint Windows Group Policy agents enable you to apply an Endpoint Windows Group Policy to a client machine and create a result session variable.

EXAMPLES

```
create endpoint-windows-group-policy { Firewall_Settings_Template }
Creates a policy for the Access Policy using the Firewall Settings template.
```

```
edit endpoint-windows-group-policy Firewall_Settings_Template
Edits the Firewall Settings Template.
```

OPTIONS

[name]
Specifies a name for the Endpoint Windows Group Policy agent.

partition
Displays the partition within which the component resides.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

policy-file
Specifies the group policy template that is applied to the client. This option is required.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-info-os, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2016-03-14 apm policy agent endpoint-windows-group-policy(1)

apm policy agent endpoint-windows-info-os

NAME

endpoint-windows-info-os - Manages an Endpoint Windows Information Operating System agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-windows-info-os component within the policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-windows-info-os [name]

modify endpoint-windows-info-os [name]

options:

app-service [[string] | none]

edit endpoint-windows-info-os [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list endpoint-windows-info-os

list endpoint-windows-info-os [[[name] | [glob] | [regex]] ...]

show running-config endpoint-windows-info-os

show running-config endpoint-windows-info-os [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show endpoint-windows-info-os

show endpoint-windows-info-os [name]

DELETE

delete endpoint-windows-info-os [name]

DESCRIPTION

You can use the endpoint-windows-info-os component to create and manage an agent that retrieves information about the Microsoft Windows operating system from the client, such as version and hotfix number.

EXAMPLES

```
create endpoint-windows-info-os MyEndpointWIOSagent { }
```

Creates the Endpoint Windows Operating System Information agent named MyEndpointWIOSagent.

```
list endpoint-windows-info-os
```

Displays a list of Endpoint Windows Operating System Information agents.

```
delete endpoint-windows-info-os MyEndpointWIOSagent delete
```

Deletes the Endpoint Windows Operating System Information agent named MyEndpointWCRagent.

OPTIONS

[name]

Specifies the name of an Endpoint Windows Info OS agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-group-policy, endpoint-machine-info, endpoint-windows-protected-workspace

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-05-30 apm policy agent endpoint-windows-info-os(1)

apm policy agent endpoint-windows-protected-workspace

NAME

endpoint-windows-protected-workspace - Manages an Endpoint Windows Protected Workspace agent.

MODULE

apm policy agent

SYNTAX

Configure the endpoint-windows-protected-workspace component within the policy agent module using the following syntax.

CREATE/MODIFY

create endpoint-windows-protected-workspace [name]

modify endpoint-windows-protected-workspace [name]

options:

allow-burn-cid [true | false]

allow-printer-use [true | false]

allow-user-switch [true | false]

allowed-network-shares [add | delete | modify | replace-all-with] {
[[string]]

}

app-service [[string] | none]

close-google-desktop-search [true | false]

usb-flash-access [all | ironkey | none]

edit endpoint-windows-protected-workspace [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list endpoint-windows-protected-workspace

list endpoint-windows-protected-workspace [[[name] | [glob] | [regex]] ...]

show running-config endpoint-windows-protected-workspace

show running-config endpoint-windows-protected-workspace [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show endpoint-windows-protected-workspace

show endpoint-windows-protected-workspace [name]

DELETE

delete endpoint-windows-protected-workspace [name]

DESCRIPTION

You can use the endpoint-windows-protected-workspace component to create and manage an agent that enables an administrator to impose limitations on applications running on Windows client machines.

OPTIONS

allow-burn-cid

Specifies that the user can burn CDs from within protected workspace. The default is false.

allow-printer-use

Specifies whether a user can print inside a protected workspace. The default is true.

allow-user-switch

Specifies whether a user can temporarily switch from a protected workspace. The default is true.

allowed-network-shares

Specifies a list of Windows network shares to which user has Write access. The default is none.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

[name]

Specifies the name of the Endpoint Windows Protected Workspace agent. This option is required.

partition

Displays the partition within which the component resides.

usb-flash-access

Specifies whether a user has access to a USB port. The default is false.

SEE ALSO

endpoint-check-software, endpoint-linux-check-file, endpoint-linux-check-process, endpoint-mac-check-file, endpoint-mac-check-process, endpoint-windows-browser-cache-cleaner, endpoint-windows-check-file, endpoint-check-machine-cert, endpoint-windows-check-process, endpoint-windows-check-registry, endpoint-windows-info-os, endpoint-machine-info

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2014-04apm policy agent endpoint-windows-protected-workspace(1)

apm policy agent external-logon-page

NAME

external-logon-page - Manages an External Logon Page agent.

MODULE

apm policy agent

SYNTAX

Configure the external-logon-page component within the policy agent module using the following syntax.

CREATE/MODIFY

```
create external-logon-page [name]
modify external-logon-page [name]
options:
  app-service [[string] | none]
  split-username [true | false]
  uri [[string]> | none]
```

```
edit external-logon-page [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list external-logon-page
list external-logon-page [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition
```

DELETE

```
delete external-logon-page [name]
```

DESCRIPTION

The External Logon Page agent creates an external Logon page that redirects the client browser.

EXAMPLES

```
create external-logon-page MyExternalLogonPageAgent { uri "MyExternalLogonPageServerURI" }
Creates the External Logon Page agent named MyExternalLogonPageAgent that is associated with the URI
MyExternalLogonPageServerURI.
```

```
create external-logon-page MyExternalLogonPageAgent { uri "%{session.my_server_uri}" }
Creates the External Logon Page agent named MyExternalLogonPageAgent with a URI of session.my_server_uri.
```

```
list external-logon-page
Displays a list of External Logon Page agents.
```

```
delete external-logon-page MyExternalLogonPageAgent
Deletes the External Logon Page agent named MyExternalLogonPageAgent.
```

OPTIONS

[name]
Specifies the name of an External Logon Page agent. This option is required.

partition
Displays the partition within which the component resides.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

split-username
Specifies whether user's input is split into username and domain. This option supports UPN style logon ID (userid@domainid) and Windows Domain User account ID (domainid\userid). The default is false. Set this to true when you want to store the username and domain separately.

uri Specifies a predefined configuration that contains several settings that you want the agent to use to configure an External Logon page. This option is required.

SEE ALSO
logon-page

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy agent external-logon-page(1)

apm policy agent http-header-modify

NAME

http-header-modify - HTTP header and cookie manipulation agent for per-request access policy.

MODULE

apm policy agent

SYNTAX

Manipulate HTTP headers or cookies within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

```
create http-header-modify [name]
```

```
modify http-header-modify [name]
```

options:

```
app-service [[string] | none]
```

```
cookie-entries [add | delete | modify | none | replace-all-with] {  
  [name] {
```

```
options:
```

```
app-service [[string] | none]
```

```
cookie-name [string]
```

```
cookie-operation [cookie-delete | cookie-update]
```

```
cookie-value [string]
```

```
}
```

```
}
```

```
header-entries [add | delete | modify | none | replace-all-with] {  
  [name] {
```

```
options:
```

```
app-service [[string] | none]
```

```
header-delimiter [string]
```

```
header-name [string]
```

```
header-operation [header-append | header-insert | header-remove | header-replace]
```

```
header-value [string]
```

```
}
```

```
}
```

```
edit http-header-modify [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list http-header-modify
```

```
list http-header-modify [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config http-header-modify
```

```
show running-config http-header-modify [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
app-service
```

```
cookie-entries
```

```
header-entries
```

```
non-default-properties
```

```
partition
```

DELETE

```
delete http-header-modify [name]
```

DESCRIPTION

You can use the http-header-modify component to create and manage a http-header-modify agent that manipulates the HTTP and Cookie headers. Operations supported for HTTP header include insert, append, replace and remove while for cookie only update and delete operations are available. Please note that this agent applies only to per-request access policy.

EXAMPLES

```
create http-header-modify MyProfile_act_http-header-modify_ag {  
  cookie-entries {
```

```

0 {
cookie-name PHPSESSID
cookie-value 1234
}
1 {
cookie-name mySession
cookie-operation cookie-delete
cookie-value 5678
}
}
header-entries {
0 {
header-name Cache-Control
header-value no-cache
}
1 {
header-delimiter ;
header-name User-Agent
header-operation header-append
header-value "Mozilla/5.0"
}
2 {
header-name Pragma
header-operation header-replace
header-value no-store
}
3 {
header-name Pragma
header-operation header-remove
header-value no-cache
}
}
}

```

In above example, http-header-modify agent named MyProfile_act_http-header-modify_ag in partition Common and adds 2 cookie and 4 HTTP header entries. cookie entry 0 updates cookie value PHPSESSID to '1234' while entry 1 deletes when cookie value 'mySession=5678'. Header entries refer to various header operations. Entry 0 inserts header 'Cache-Control: 5678'. Entry 1 updates value of header User-Agent to include ';Mozilla/5.0'. Entry 2 replaces value of header Pragma by 'no-store'. Finally, 3 entry will remove header-value matching 'Pragma: no-cache' from HTTP headers.

list http-header-modify

Displays a list of http-header-modify agents.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

name Specifies the name of HTTP header modify agent. This setting is required.

cookie-entries

Specifies a list of entries specifying cookie header manipulations.

cookie-name

Specifies cookie name to match. This setting is required for cookie-entries.

cookie-operation

Specifies operation on the cookie name specified in cookie-name. Possible values include cookie-update and cookie-delete.

cookie-update

Update the cookie value in the cookie-name, cookie-value pair. This is the default cookie-operation.

cookie-delete

Delete the cookie that matches the cookie-name, cookie-value pair.

cookie-value

Specifies cookie value to be operated on. This is required for cookie-entries.

header-entries

Specifies a list of entries specifying HTTP header manipulations.

header-delimiter

Specifies delimiter character to use when header-operation is header-append.

header-name

Specifies HTTP header to match to be operated on. This setting is required for header-entries.

header-operation

Specifies operation on the HTTP header specified in header-name. Options include the following:

header-append

Append header-value to the value of HTTP header header-name delimited by header-delimiter.

header-insert

Insert HTTP header containing header-name, header-value pair. This is the default header-operation.

header-remove

Remove the HTTP header, value pair that matches the header-name, header-value pair.

header-replace

Replace value of HTTP header matching header-name by value header-value.

header-value

Specifies HTTP header value to be operated on. This is required for header-entries.

partition

Displays the partition within which the component resides.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm policy agent http-header-modify(1)

apm policy agent ip-geolocation-lookup

NAME

ip-geolocation-lookup - Manages an IP Geolocation Lookup agent.

MODULE

apm policy agent

SYNTAX

Configure the ip-geolocation-lookup component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ip-geolocation-lookup [name]
modify ip-geolocation_lookup [name]
options:
  app-service [[string] | none]
```

```
edit ip_geolocation_lookup [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list ip_geolocation_lookup
list ip_geolocation_lookup [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition
```

DELETE

```
delete ip_geolocation_lookup [name]
```

DESCRIPTION

You can use the ip_geolocation_lookup component to create and manage an IP Geolocation Lookup agent.

EXAMPLES

```
create ip_geolocation_lookup example_ip_geolocation_lookup_ag
Creates the example_ip_geolocation_lookup_ag IP Geolocation Lookup agent that searches for the IP Geolocation of the chosen IP address.
```

```
delete ip_geolocation_lookup example_ip_geolocation_lookup_ag delete
Deletes the IP Geolocation Lookup agent named example_ip_geolocation_lookup_ag.
```

OPTIONS

[name]
Specifies the name of a IP Geolocation Lookup agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017-2018. All rights reserved.

BIG-IP 2018-04-23 apm policy agent ip-geolocation-lookup(1)

apm policy agent ip-reputation-lookup

NAME

ip-reputation-lookup - Manages an IP Reputation Lookup agent.

MODULE

apm policy agent

SYNTAX

Configure the ip-reputation-lookup component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create ip-reputation-lookup [name]

modify ip-reputation_lookup [name]

options:

app-service [[string] | none]

edit ip_reputation_lookup [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ip_reputation_lookup

list ip_reputation_lookup [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

DELETE

delete ip_reputation_lookup [name]

DESCRIPTION

You can use the ip_reputation_lookup component to create and manage an IP Reputation Lookup agent.

EXAMPLES

create ip_reputation_lookup example_ip_reputation_lookup_ag

Creates the example_ip_reputation_lookup_ag IP Reputation Lookup agent that searches for the IP Reputation of the chosen IP address.

delete ip_reputation_lookup example_ip_reputation_lookup_ag delete

Deletes the IP Reputation Lookup agent named example_ip_reputation_lookup_ag.

OPTIONS

[name]

Specifies the name of a IP Reputation Lookup agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

apm policy agent irule-event

NAME

irule-event - Manages an iRule Event agent.

MODULE

apm policy agent

SYNTAX

Configure the irule-event component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create irule-event [name]

modify irule-event [name]

options:

app-service [[string] | none]

id [[string] | none]

expect-data [[client-accepted] | [ssl-client-hello] | [ssl-cert-available] | [http] | [none]]

edit irule-event [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list irule-event

list irule-event [[[name] | [glob] | [regex]] ...]

show running-config irule-event

show running-config irule-event [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show irule-event

show irule-event [name]

DELETE

delete irule-event [name]

DESCRIPTION

You can use the irule-event component to add a custom Access iRule event to an access policy. This agent enables you to combine access policy execution with iRule execution.

For example, you can retrieve the current agent ID (using an iRule command `ACCESS::policy agent_id`) to determine which of the iRule agents raised the event and then perform some custom logic execution.

The Per-Request Policy's Access iRule event functions similarly to the Access iRule event in the main access policy, but retrieval of the current agent ID uses the iRule command `"ACCESS::perflow get perflow.irule_agent_id"` to determine which of the iRule agents raised the event.

The expect-data component is used in the Per-Request Policy when adding an Access iRule event. You must specify the last piece of data that the policy should wait for before raising the event. The default is HTTP.

EXAMPLES

```
when ACCESS_POLICY_AGENT_EVENT { if {[ACCESS::policy agent_id] eq "lastLogon" } { # our limit in seconds set
2weeks 1209600 # diff in 100 nanosecond increments between MS time attribute (year 1601) and start of epoch
set offset 11644473600000 set adtime "[ACCESS::session data get session.ad.last.attr.lastLogon]" # convert
adtime to milliseconds set millisecs [expr {$adtime / 10000}] # subtract offset set lastlogintime [expr
{$millisecs - $offset}] # convert to seconds because milliseconds for 'now' were negative (maybe vmware issue)
set secs [expr {$lastlogintime / 1000}] set now [clock seconds] # finally calculate the difference set diff
[expr {$now - $secs}] log local0. "lastLogon: $diff seconds from current time" if { $diff $2weeks } {
ACCESS::session data set session.custom.lastLogonWithin2Weeks 0 } else { ACCESS::session data set
session.custom.lastLogonWithin2Weeks 1 } } }
```

In this example, `ACCESS_POLICY_AGENT_EVENT` gathers data containing the users whose last logon was within the last two weeks. Note that you can access session variables and create new session variables inside this event.

```
list irule-event all
```

Displays a list of OAM agents.

```
delete irule-event my_irule_agent
```

Deletes the iRule Event agent named `my_irule_agent`.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

id Specifies the ID of the iRule event. The default is none. You can use the ID to determine which agent caused the ACCESS_POLICY_AGENT_EVENT. You can also use the ID to perform different processing inside iRule for different agents.

[name]

Specifies the name of the component. This option is required.

partition

Displays the partition within which the component resides.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2017-11-28 apm policy agent irule-event(1)

apm policy agent kerberos

NAME

kerberos - Manages a Kerberos agent.

MODULE

apm policy agent

SYNTAX

Configure the kerberos component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create kerberos [name]

modify kerberos [name]

options:

app-service [[string] | none]

max-logon-attempt [integer]

server [string]

edit kerberos [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list kerberos

list kerberos [[[name] | [glob] | [regex]] ...]

show running-config kerberos

show running-config kerberos [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show kerberos

show kerberos [name]

DELETE

delete kerberos [name]

DESCRIPTION

You can use the kerberos component to create and manage a Kerberos agent.

EXAMPLES

create kerberos my_kerberos_agent

Creates a Kerberos agent named my_kerberos_agent.

list kerberos all

Displays a list of Kerberos agents.

delete kerberos my_kerberos_agent

Deletes the Kerberos agent named `my_kerberos_agent`.

OPTIONS

`[name]`

Specifies the name of the component. This option is required.

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`max-logon-attempt`

Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

`partition`

Displays the partition within which the component resides.

`server`

Specifies the name of the Kerberos server. This option is required.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy agent kerberos(1)

apm policy agent l7-protocol-lookup

NAME

`l7-protocol-lookup` - Manages a L7 Protocol Lookup agent.

MODULE

apm policy agent

SYNTAX

Configure the `l7-protocol-lookup` component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

```
create l7-protocol-lookup [name]
modify l7_protocol_lookup [name]
options:
  app-service [[string] | none]
```

```
edit l7_protocol_lookup [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list l7_protocol_lookup
list l7_protocol_lookup [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition
```

DELETE

```
delete l7_protocol_lookup [name]
```

DESCRIPTION

You can use the `l7_protocol_lookup` component to create and manage a L7 Protocol Lookup agent.

EXAMPLES

```
create l7_protocol_lookup example_l7_protocol_lookup_ag
Creates the example_l7_protocol_lookup_ag L7 Protocol Lookup agent that identifies the layer 7 protocol.
```

```
delete l7_protocol_lookup example_l7_protocol_lookup_ag delete
Deletes the L7 Protocol Lookup agent named example_l7_protocol_lookup_ag.
```

OPTIONS

[name]

Specifies the name of a L7 Protocol Lookup agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017-2018. All rights reserved.

BIG-IP 2018-04-21 apm policy agent l7-protocol-lookup(1)

apm policy agent logging

NAME

logging - Manages a Logging agent.

MODULE

apm policy agent

SYNTAX

Configure the logging component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create logging [name]

modify logging [name]

options:

app-service [[string] | none]

log-message [[string] | none]

variables [[string] | none]

edit logging [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list logging

list logging [[[name] | [glob] | [regex]] ...]

show running-config logging

show running-config logging [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

log-message

non-default-properties

partition

variables

DELETE

delete logging [name]

DESCRIPTION

You can use the logging component to create and manage a logging agent that monitors the value of session variables and identifies the path taken by access policy execution. A logging agent can also be used to create and monitor custom or predefined session variables. Note that a session variable may or may not exist depending on the result of the access policy execution.

EXAMPLES

```
create logging MyProfile_act_logging_ag {
```

```
  variables
```

```
  {
```

```
{session-var "session.logon.*"}
```

```
{session-var "session.windows_check_file.Company8profile_act_file_check_ag.item_x.filename"}
```

```
  }
```

```
}
```

Creates the logging agent named MyProfile_act_logging_ag in partition Common and adds two session variables that define actions that the agent logs: session.logon.* indicates to log application logon attempts and

session.windows_check.file.Company8profile_act_file_check_ag.item_x.filename indicates to log the outcome of the file check on the client. The x in item_x indicates the order of the files in the list configured for the file checker. The list starts with index 0 (zero).

```
create logging MyProfile_act_logging_ag {
  log-message "Logon Name: %{session.logon.last.username}."
}
```

Above example applies only to logging agent tied to per-request access policy. Here logging agent named MyProfile_act_logging_ag in partition Common will print log messages containing logon name. This removes the requirement to configure variables separately for per-request access policy.

list logging

Displays a list of logging agents.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

log-message

Specifies the log message to display. This option is required. For per-request access policy only, this option can contain session or per-request variables. However, session variables containing wildcard (*) are not supported.

name Specifies the name of a logging agent. This option is required.

partition

Displays the partition within which the component resides.

variables

Adds a variable to or deletes a variable from a logging agent. You use the sessionvar option to specify a session variable that indicates which actions the system logs. This option does not apply to per-request access policy agent.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2016. All rights reserved.

BIG-IP 2016-03-14 apm policy agent logging(1)

apm policy agent logon-page

NAME

logon-page - Manages a Logon Page agent.

MODULE

apm policy agent

SYNTAX

Configure the logon-page component within the policy agent module using the following syntax.

CREATE/MODIFY

create logon-page [name]

modify logon-page [name]

options:

app-service [[string] | none]

basic-auth-realm [[string] | none]

customization-group [[string] | none]

field-modifiable1 [true | false]

field-modifiable2 [true | false]

field-modifiable3 [true | false]

field-modifiable4 [true | false]

field-modifiable5 [true | false]

field-type1 [checkbox | none | password | text]

field-type2 [checkbox | none | password | text]

field-type3 [checkbox | none | password | text]

field-type4 [checkbox | none | password | text]

field-type5 [checkbox | none | password | text]

http-401-auth-level [basic | basic-negotiate | negotiate | none]

post-var-name1 [[integer] | none]

post-var-name2 [[integer] | none]

post-var-name3 [[integer] | none]

post-var-name4 [[integer] | none]

post-var-name5 [[integer] | none]
session-var-name1 [[integer] | none]
session-var-name2 [[integer] | none]
session-var-name3 [[integer] | none]
session-var-name4 [[integer] | none]
session-var-name5 [[integer] | none]
clean-sess-var1 [true | false]
clean-sess-var2 [true | false]
clean-sess-var3 [true | false]
clean-sess-var4 [true | false]
clean-sess-var5 [true | false]
split-username [true | false]
type [401 | form-based]

edit logon-page [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list logon-page

list logon-page [[[name] | [glob] | [regex]] ...]

show running-config logon-page

show running-config logon-page [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
partition

show logon-page

show logon-page [name]

DELETE

delete logon-page [name]

DESCRIPTION

You can use the logon-page component to create and manage a Logon Page agent. This agent creates a logon page that includes the form in which users input the credentials required by an access policy. You can use the customization-group option to customize the logon page.

EXAMPLES

```
create logon-page MyLogonPageAgent my { type 401 basic-auth-realm myrealm split-username false  
http-401-auth-level none }
```

Creates a basic authentication Logon Page agent named MyLogonPageAgent that results in a 401 response.

```
list logon-page
```

Displays a list of Logon Page agents.

```
delete logon-page MyLogonPageAgent
```

Deletes the Logon Page agent named MyLogonPageAgent.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

basic-auth-realm

Specifies the system being accessed for HTTP basic authentication. This value is shown in the 401 response. Use this option only for basic authentication Logon pages.

clean-sess-var1 - clean-sess-var5

Specifies whether session variable under corresponding sess-var-name would be cleaned before logon page appearance (i.e. agent execution) or not. Works only with form-based authentication.

customization-group

Specifies a predefined configuration that contains several settings that you want the agent to use to configure a logon page. This setting is required, and the customization group that you assign must be of the type logon. Use this option only for basic authentication Logon pages.

field-modifiable1 - field-modifiable5

Specifies whether the user can modify the contents of the field on a form-based Logon page. The default is true. You can use this option to display read-only information. A Logon page contains can have a maximum of five fields. Use this option only for form-based Logon pages.

field-type1 - field-type5

Specifies the type of fields on a form-based Logon page. The default is text. Use this option only for form-based Logon pages. The options are:

checkbox

none

password

text

http-401-auth-level

Use this option only for basic authentication Logon pages. The options are:

basic

basic-negotiate
negotiate
none

[name]

Specifies the name of a Logon Page agent. This setting is required.

partition
Displays the partition within which the component resides.

post-var-name1 - post-var-name5
Specifies the name of the variable that is sent with POST request. Use this option only for form-based Logon pages.

sess-var-name1 - sess-var-name5
Specifies the session variable from which the initial value is taken. Use this option only for form-based Logon pages.

split-username
Specifies whether the user's input is split into username and domain. This option supports UPN style logon ID (userid@domainid) and Windows Domain User account ID (domainid\userid). The default is false. Set this to true when you want to store the username and domain separately.

Use this option only for basic authentication Logon pages.

type Specifies the type of logon page that appears. The options are:

401 Displays a basic HTTP authentication form.

form-based
Displays a logon page.

SEE ALSO
external-logon-page

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2016-10-03 apm policy agent logon-page(1)

apm policy agent message-box

NAME
message-box - Manages a Message Box agent.

MODULE
apm policy agent

SYNTAX
Configure the message-box component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY
create message-box [name]
modify message-box [name]
options:
app-service [[string] | none]
customization-group [string]

edit message-box [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list message-box
list message-box [[[name] | [glob] | [regex]] ...]
show running-config message-box
show running-config message-box [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
partition

show message-box
show message-box [name]

DELETE
delete message-box [name]

DESCRIPTION

You can use the message-box agent to create, display, or delete a Message Box agent. You cannot use the command line interface to create or modify the messages that display in a message box. You can also edit customizable messages using the visual policy editor.

EXAMPLES

```
create message-box MyMessageBoxAgent { customization group "MyMessageBoxCG" }  
Creates the Message Box agent named MyMessageBoxAgent that is associated with the customization group named MyMessageBoxCG.
```

```
list message-box  
Displays a list of Message Box agents.
```

```
delete message-box MyMessageBoxAgent  
Deletes the Message Box agent named MyMessage BoxAgent.
```

OPTIONS

[name]
Specifies the name of a Message Box agent. This option is required.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

customization-group
Specifies the name of the customization group that contains the messages you want to apply to an access policy. This option is required.

partition
Displays the partition within which the component resides.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-12-20 apm policy agent message-box(1)

apm policy agent oam

NAME

oam - Manages an OAM agent.

MODULE

apm policy agent

SYNTAX

Warning: This page is obsolete. AAA OAM agent is no longer supported.

Configure the oam component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

```
create oam [name]  
modify oam [name]  
options:  
  app-service [[string] | none]  
  max-logon-attempt [integer]  
  server [[string] | none]  
  show-extended-error [true | false]  
  url [[string] | none]
```

```
edit oam [ [ [name] | [glob] | [regex] ] ... ]  
options:  
  all-properties  
  non-default-properties
```

DISPLAY

```
list oam  
list oam [ [ [name] | [glob] | [regex] ] ... ]  
show running-config oam  
show running-config oam [ [ [name] | [glob] | [regex] ] ... ]
```

options:
all-properties
non-default-properties
partition

show oam
show oam [name]

DELETE
delete oam [name]

DESCRIPTION

You can use the oam component to create and manage an OAM agent.

EXAMPLES

```
create oam oam_agent1 { server oam10g max-logon-attempt 3 show-extended-error false url  
"http://www.mydomain.com/protected/" }
```

Creates an OAM agent named oam_agent1 that uses authentication server oam10g and prompts a user for credentials three times before denying access to http://www.mydomain.com/protected/.

```
modify oam oam_agent1 max-logon-attempt 4  
list oam all
```

Displays a list of OAM agents.

```
delete oam my_oam_agent
```

Deletes the OAM agent named my_tacacsplus_agent.

OPTIONS

[name]

Specifies the name of the component. This option is required.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

max-logon-attempt

Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

partition

Displays the partition within which the component resides.

server

Specifies the name of the OAM server used for user authentication. This option is required.

url Specifies the URL of the resource that is protected by the OAM server. It is used to authenticate the user using the specified user credentials. This option is required, and you must specify an absolute URL.

An absolute URL specifies the exact location of a file or directory on the Internet.

show-extended-error

Specifies to display a verbose error message on the retry logon page. The default value is false.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm policy agent oam(1)

apm policy agent oauth-authz

NAME

oauth-authz - Manages an OAuth Authorization agent.

MODULE

apm policy agent

SYNTAX

Configure the oauth-authz component within the policy agent module using the following syntax.

CREATE/MODIFY

create oauth-authz [name]

```

modify oauth-authz [name]
  options
    app-service [[string] | none]
    audience ( | none) [add | delete]
    customization-group [[string] | none]
    entries [ add | delete | modify | none | replace-all-with] {
[entry-name] {
  options:
    app-service [[string] | none]
    expression [[string] | none]
    id-token-claim-entries [add | delete | modify | none | replace-all-with] {
[id-token-claim-entry-name] {
  options:
    app-service [[string] | none]
    claim-name [claim-name]
    claim-value [[string] | none]
}
}
  jwt-access-token-claim-entries [add | delete | modify | none | replace-all-with] {
[jwt-access-token-claim-entry-name] {
  options:
    app-service [[string] | none]
    claim-name [claim-name]
    claim-value [[string] | none]
}
}
  scope-entries [add | delete | modify | none | replace-all-with] {
[scope-entry-name] {
  options:
    app-service [[string] | none]
    scope-name [scope-name]
    scope-value [[string] | none]
}
}
  userinfo-claim-entries [add | delete | modify | none | replace-all-with] {
[userinfo-claim-entry-name] {
  options:
    app-service [[string] | none]
    claim-name [claim-name]
    claim-value [[string] | none]
}
}
}
}
  prompt-for-authorization [true | false]
  subject [[string] | none]

```

```

edit oauth-authz [ [ [name] | [glob] | [regex] ] ... ]
  options:
    all-properties
    non-default-properties

```

DISPLAY

```

list oauth-authz
list oauth-authz [ [ [name] | [glob] | [regex] ] ... ]
show running-config oauth-authz
show running-config oauth-authz [ [ [name] | [glob] | [regex] ] ... ]
  options:
    all
    all-properties
    app-service
    customization-group
    entries
    non-default-properties
    one-line
    partition
    prompt-for-authorization
    recursive

```

DELETE

```

delete oauth-authz [name]

```

DESCRIPTION

You can use the `oauth-authz` component to create and manage an OAuth Authorization agent that provides OAuth Authorization server functionality, and also manage scopes and claims to provide different level of access control based on end user's role or any other criteria. For JWT type tokens, you can use the agent to manage audience and subject values.

EXAMPLES

```

create oauth-authz MyOAuthAuthzAgent {
  audience add { "company-oauth-rs.com" "partner-oauth-rs.com" }
  customization-group "company_authz"
  entries add {
  0 {
    expression "expr {return true}"
    id-token-claim-entries add {
    0 {

```

```

claim-name "group"
claim-value "%{session.ad.last.attr.memberOf}"
}
}
jwt-access-token-claim-entries add {
0 {
claim-name "group"
claim-value "%{session.ad.last.attr.memberOf}"
}
1 {
claim-name "profile"
claim-value "https://company.com/username"
}
}
userinfo-claim-entries add {
0 {
claim-name "profile"
claim-value "https://company.com/username"
}
}
}
scope-entries add {
0 {
scope-name "name"
scope-value "%{session.logon.last.name}"
}
1 {
scope-name "email"
scope-value "test@company.com"
}
2 {
scope-name "domain"
scope-value "%{session.logon.last.domain}"
}
}
}
1 {
expression "expr {[mcget {session.logon.last.name}] == "testuser"}"
jwt-access-token-claim-entries add {
0 {
claim-name "service"
claim-value "medium"
}
}
scope-entries add {
0 {
scope-name "project"
scope-value "project-one"
}
}
}
}
}
subject "%{session.assigned.uuid}"
}
}

```

Creates an OAuth Authorization agent named MyOAuthAuthzAgent that uses customization group company_authz to customize the OAuth Authorization page.

The agent associates these scopes name, email, and domain and their values to each access token because the first expression always evaluates to true. If an id_token is issued, it contains claim group. If the token type issued is JWT, each access token also includes claims group and profile and the UserInfo response will contain claim profile. The agent also associates scope project with value project-one to the token if the user is testuser. If JWT access token is issued to user testuser, it will contain claim service with value medium.

The agent additionally includes audience containing values company-oauth-rs.com and partner-oauth-rs.com along with subject and its value if the token issued is a JWT access token.

```
list oauth-authz
```

Displays a list of OAuth Authorization agents.

```
delete oauth-authz MyOAuthAuthzAgent
```

Deletes the OAuth Authorization agent named MyOAuthAuthzAgent.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

audience

Specifies a list of audience values used in JWT tokens issued. If audience list is specified in the OAuth Authorization agent, it overwrites the values in OAuth profile and OAuth client app.

customization-group

Specifies the customization group that defines the appearance of the OAuth Authorization page.

entries

Specifies a list of entries consisting of an expression and a list of scope entries. If the expression evaluates to true, then the OAuth Authorization agent associates the corresponding list of scope entries to an issued token. Scope entries determine the access control that the OAuth Authorization server requests on behalf of the client application.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

expression

Specifies the expression that you want an OAuth Authorization agent to use to verify in order to associate the corresponding scopes to an issued token. You can use the following operators: AND, OR, NOT, (and).

id-token-claim-entries

Specifies a list of entries consisting of an ID token claim name and its value.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

claim-name

Specifies the name of the claim.

claim-value

Specifies a value to the corresponding claim. This value can be any string or session variable.

jwt-access-token-claim-entries

Specifies a list of entries consisting of a JWT access-token claim name and its value.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

claim-name

Specifies the name of the claim.

claim-value

Specifies a value to the corresponding claim. This value can be any string or session variable.

scope-entries

Specifies a list of entries consisting of a scope name and its value.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

scope-name

Specifies the name of the scope.

scope-value

Specifies a value to the corresponding scope. This value can be any string or session variable.

userinfo-claim-entries

Specifies a list of entries consisting of a UserInfo claim name and its value.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

claim-name

Specifies the name of the claim.

claim-value

Specifies a value to the corresponding claim. This value can be any string or session variable.

[name]

Specifies the name of the OAuth Authorization agent. This setting is required.

partition

Displays the partition within which the component resides.

prompt-for-authorization

Specifies whether the OAuth Authorization page, for user authorization, is displayed. This is applicable

only for "Authorization code" and "Implicit" grants.

subject

Specifies the value of subject in JWT tokens issued. If subject is specified in the OAuth Authorization agent, it overwrites the the value specified in the OAuth profile.

SEE ALSO

apm oauth oauth-scope apm oauth oauth-claim

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2017. All rights reserved.

BIG-IP 2017-10-20 apm policy agent oauth-authz(1)

apm policy agent request-classification

NAME

request-classification - Manages a Request Classification agent.

MODULE

apm policy agent

SYNTAX

Configure the request-classification component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create request-classification [name]

modify request-classification [name]

options:

app-service [[string] | none]

edit api-server-selection [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list request-classification

list request-classification [[[name] | [glob] | [regex]] ...]

show running-config request-classification

show running-config request-classification [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

DELETE

delete request-classification [name]

DESCRIPTION

Use the request-classification component to create, modify, display, or delete a Request Classification agent.

A Request Classification agent classifies the API requests based on their URIs and http methods according to the API path definition in the corresponding API Protection profile.

EXAMPLES

create request-classification my_rca_ag

Creates the my_rca_ag Request Classification agent.

list request-classification

Displays a list of Request Classification agents.

delete request-classification my_rca_ag

Deletes the my_rca_ag Request Classification agent.

OPTIONS

[name]

Specifies the name of a Request Classification agent. This option is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot

modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

api-protection profile apiprotection

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-07-12 apm policy agent request-classification(1)

apm policy agent resource-assign

NAME

resource-assign - Manages a Resource Assign agent.

MODULE

apm policy agent

SYNTAX

Configure the resource-assign component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create resource-assign [name]

modify resource-assign [name]

options:

app-service [(string) | none]

rules (| none)

type [acls | general | resources | webtop-and-webtop-links]

edit resource-assign [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list resource-assign

list resource-assign [[[name] | [glob] | [regex]] ...]

show running-config resource-assign

show running-config resource-assign [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show resource-assign

show resource-assign [name]

DELETE

delete resource-assign [name]

DESCRIPTION

You can use the resource-assign component to create and manage an agent that assigns an ACL, a resource group, or both to an access policy. A resource group is a collection of resources, access control lists, and protection criteria, which includes your company intranet servers, applications, and network shares. An ACL is a set of restrictions associated with a resource or favorite that defines access for users and groups.

EXAMPLES

```
create resource-assign MyAssignResourceAgent my rules { { expression "expr { [mcget {session.ad.last.authresult}] == 1 }" webtop-links add { google } } }
```

Creates the Resource Assign agent named MyAssignResourceAgent and assigns webtop-link google when authentication is passed.

```
list resource-assign all
```

Displays a list of Resource Assign agents.

```
delete resource-assign MyAssignResourceAgent
```

Deletes the Resource Assign agent named MyAssignResourceAgent.

OPTIONS

[name]

Specifies the name of the Resource Assign agent. This option is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

rules

Adds a rule to or deletes a rule from the Resource Assign agent. You can use the following attributes to define a rule:

acl Specifies an access control list that this rule assigns to users.

connectivity-resource-group

Specifies the name of the connectivity resource group to which this rule applies.

expression

Specifies the expression that indicates which resource groups this rule assigns to users.

type Specifies the type of Resource Assign agent. The default is general.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy agent resource-assign(1)

apm policy agent response-selection

NAME

response-selection - Manages a Response Selection agent.

MODULE

apm policy agent

SYNTAX

Configure the response-selection component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create response-selection [name]

modify response-selection [name]

options:

app-service [[string] | none]

response [[string] | none]

edit response-selection [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list response-selection

list response-selection [[[name] | [glob] | [regex]] ...]

show running-config response-selection

show running-config response-selection [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show response-selection

show response-selection [name]

DELETE

delete response-selection [name]

DESCRIPTION

You can use the response-selection component to create a Response Selection agent to assign a response config object for error response generation.

EXAMPLES

create response-selection my_resp_sel_ag response my_response

Creates the my_resp_sel_ag Response Selection agent.

list response-selection

Displays a list of Response Selection agents.

```
delete response-selection my_resp_sel_agent
```

Deletes the Response Selection agent named my_resp_sel_agent.

OPTIONS

`[name]`
Specifies the name of a Response Selection agent. This option is required.

`partition`
Displays the partition within which the component resides.

`app-service`
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`response`
Specifies the response config object name.

SEE ALSO

api-protection response

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-07-12 apm policy agent response-selection(1)

apm policy agent route-domain-selection

NAME

route-domain-selection - Manages a Route Domain Selection agent.

MODULE

apm policy agent

SYNTAX

Configure the route-domain-selection component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

```
create route-domain-selection [name]
```

```
modify route-domain-selection [name]
```

options:

```
app-service [[string] | none]
```

```
location-specific [true | false]
```

```
route-domain [[integer] | none]
```

```
snat [automap | none]
```

```
snatpool [[string] | none]
```

```
edit route-domain-selection [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list route-domain-selection
```

```
list route-domain-selection [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config route-domain-selection
```

```
show running-config route-domain-selection [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
partition
```

```
show route-domain-selection
```

```
show route-domain-selection [name]
```

DELETE

```
delete route-domain-selection [name]
```

DESCRIPTION

You can use the route-domain-selection component to create a Route Domain Selection agent.

EXAMPLES

create route-domain-selection my_rds_ag route-domain 0 snat automap
Creates the my_rds_ag Route Domain Selection agent.

show route-domain-selection
Displays a list of Route Domain Selection agents.

delete route-domain-selection my_rd_selection_agent
Deletes the Route Domain Selection agent named my_rd_selection_agent.

OPTIONS

[name]
Specifies the name of a Variable Assignment agent. This option is required.

partition
Displays the partition within which the component resides.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

location-specific
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

route-domain
Specifies the route domain. The default is 0 (zero).

snat
automap
none Snat is not used.

snatpool

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-24 apm policy agent route-domain-selection(1)

apm policy agent server-cert-response-control

NAME

server-cert-response-control - Manages a Server Cert Response Control agent.

MODULE

apm policy agent

SYNTAX

Configure the server-cert-response-control component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create server-cert-response-control [name]
modify server-cert-response-control [name]

options:

app-service [[string] | none]
action [integer]

edit server-cert-response-control [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list server-cert-response-control

list server-cert-response-control [[[name] | [glob] | [regex]] ...]

options:

all-properties
app-service
non-default-properties
partition

DELETE
delete server-cert-response-control [name]

DESCRIPTION
You can use the server-cert-response-control component to create and manage a Server Cert Response Control agent.

EXAMPLES
create server-cert-response-control example_server_cert_response_control_ag
Creates the example_server_cert_response_control_ag Server Cert Response Control agent that allows admin to either ignore or mask expired/untrusted server certificate.

delete server-cert-response-control example_server_cert_response_control_ag delete
Deletes the Server Cert Response Control agent named example_server_cert_response_control_ag.

OPTIONS
[name]
Specifies the name of a Server Cert Response Control agent. This setting is required.

partition
Displays the partition within which the component resides.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

action
Allows admin to specify how to handle connections from a client whose server certificate is either expired or untrusted. The default is ignore which specifies that the system ignores untrusted/expired certificate and may allow the connection. When mask option is selected, end users will see the block page without the cert expired/untrusted warning message from clients/browsers.

SEE ALSO
COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018-2019. All rights reserved.

BIG-IP 2018-11-21 apm policy agent server-cert-response-control(1)

apm policy agent server-cert-status

NAME
server-cert-status - Manages a Server Cert Status agent.

MODULE
apm policy agent

SYNTAX
Configure the server-cert-status component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY
create server-cert-status [name]
modify server-cert-status [name]
options:
app-service [[string] | none]

edit server-cert-status [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list server-cert-status
list server-cert-status [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
partition

DELETE
delete server-cert-status [name]

DESCRIPTION
You can use the server-cert-status component to create and manage a Server Cert Status agent.

EXAMPLES

```
create server-cert-status example_server_cert_status_ag
```

Creates the example_server_cert_status_ag Server Cert Status agent that checks the server certificate status.

```
delete server-cert-status example_server_cert_status_ag delete
```

Deletes the Server Cert Status agent named example_server_cert_status_ag.

OPTIONS

[name]

Specifies the name of a Server Cert Status agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018-2019. All rights reserved.

BIG-IP 2018-11-19 apm policy agent server-cert-status(1)

apm policy agent session-check

NAME

session-check - Manages a Session Check agent.

MODULE

apm policy agent

SYNTAX

Configure the session-check component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

```
create session-check [name]
```

```
modify session-check [name]
```

options:

```
app-service [[string] | none]
```

```
edit session-check [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list session-check
```

```
list session-check [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
partition
```

DELETE

```
delete session-check [name]
```

DESCRIPTION

You can use the session-check component to create and manage an agent that check for the existence of access policy session in a flow.

EXAMPLES

```
create session-check example_session_check_ag
```

Creates the example_session_check_ag Session Check agent that performs the access policy session existence check for the flow.

```
delete session-check ExampleSessionCheckAgent delete
```

Deletes the Session Check agent named ExampleSessionCheckAgent.

OPTIONS

[name]

Specifies the name of a Session Check agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017-2018. All rights reserved.

BIG-IP 2018-04-21 apm policy agent session-check(1)

apm policy agent ssl-check

NAME

ssl-check - Manages a SSL Check agent.

MODULE

apm policy agent

SYNTAX

Configure the ssl-check component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create ssl-check [name]

modify ssl-check [name]

options:

app-service [[string] | none]

edit ssl-check [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ssl-check

list ssl-check [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

DELETE

delete ssl-check [name]

DESCRIPTION

You can use the ssl-check component to create and manage an agent that check for the existence of SSL traffic in a flow.

EXAMPLES

create ssl-check example_ssl_check_ag

Creates the example_ssl_check_ag SSL Check agent that performs the SSL traffic existence check for the flow.

delete ssl-check ExampleSSLCheckAgent delete

Deletes the SSL Check agent named ExampleSSLCheckAgent.

OPTIONS

[name]

Specifies the name of a SSL Check agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017-2018. All rights reserved.

BIG-IP 2018-04-21 apm policy agent ssl-check(1)

apm policy agent tacacsplus

NAME

tacacsplus - Manages a TACACS+ agent.

MODULE

apm policy agent

SYNTAX

Configure the tacacsplus component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create tacacsplus

modify tacacsplus

options:

app-service [[string] | none]

max-logon-attempt [integer]

server [[string] | none]

edit tacacsplus [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tacacsplus

list tacacsplus [[[name] | [glob] | [regex]] ...]

show running-config tacacsplus

show running-config tacacsplus [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show tacacsplus

show tacacsplus [name]

DELETE

delete tacacsplus [name]

DESCRIPTION

You can use the tacacsplus component to create and manage a TACACS+ agent.

EXAMPLES

list tacacsplus all

Displays a list of TACACS+ agents.

delete tacacsplus my_tacacsplus_agent

Deletes the TACACS+ agent named my_tacacsplus_agent.

OPTIONS

[name]

Specifies the name of the component. This option is required.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

max-logon-attempt

Specifies the maximum number of opportunities that users have to re-enter credentials after their first attempt to log in fails. If you set this value to a number from 2 to 5 inclusive, the system allows users the specified number of opportunities to log in after the first attempt to log in fails. If you set the value to 1, the system does not allow a second log in opportunity after a first log in attempt fails. The default value is 3.

partition

Displays the partition within which the component resides.

server

Specifies the name of the TACACS+ server. This option is required.

SEE ALSO
COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy agent tacacsplus(1)

apm policy agent variable-assign

NAME

variable-assign - Manages a Variable Assignment agent.

MODULE

apm policy agent

SYNTAX

Configure the variable-assign component within the policy agent module using the syntax shown in the following sections.

CREATE/MODIFY

create variable-assign [name]

modify variable-assign [name]

options:

app-service [[string] | none]

type [citrix-smart-access | general | intranet-webtop | sso-cred-mapping | virtual-keyboard]

variables { [varname [name] expression {[string]}] }

edit variable-assign [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list variable-assign

list variable-assign [[[name] | [glob] | [regex]] ...]

show running-config variable-assign

show running-config variable-assign [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

show variable-assign

show variable-assign [name]

DELETE

delete variable-assign [name]

DESCRIPTION

You can use the variable-assign component to create and manage an agent that assigns one or more variables to an access policy. F5 Networks recommends that you use the visual policy editor to create complex variable assignments.

EXAMPLES

```
create variable-assign username_variable_assign_ag { variables { varname "session.logon.last.username"
expression "[{mcget {session.ssl.cert.cn}}]" } }
```

Creates the username_variable_assign_ag Variable Assignment agent that automatically assigns the value of the common name field in the client certificate to the username field of the logon page. This is useful when an access policy contains the Variable Assignment agent between the client certification and the AAA Active Directory server query actions.

```
create variable-assign acl_variable_assign_ag { variables { varname
```

```
"config.connectivity_resource_network_access.MyprofileNR2.acl_name" expression "expr {\\"MY_ACL1\\"}" } }
```

Creates a Variable Assignment agent that carries out a configured ACL when a particular branch in the access policy is followed, using the Variable Assignment agent to populate the appropriate variables with the ACL name.

```
show variable-assign
```

Displays a list of Variable Assignment agents.

```
delete variable-assign MyAssignVariableAgent delete
```

Deletes the Variable Assignment agent named MyAssignVariableAgent.

OPTIONS

[name]

Specifies the name of a Variable Assignment agent. This setting is required.

partition

Displays the partition within which the component resides.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

type Specifies the type of agent. The default is general.

variables

Adds a variable to or deletes a variable from the Variable Assignment agent. You must specify the following attributes for each variable:

expression

A Tcl expression that the system evaluates, and then assigns the value of the expression to a specific property of the assigned Network Access resource or to a newly created session variable.

varname

A variable that forms the left-hand side of the expression. You can use the name of an existing variable or a new session variable.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-12-20 apm policy agent variable-assign(1)

apm policy customization-group

NAME

customization-group - Manages a customization group.

MODULE

apm policy

SYNTAX

Warning: F5 Networks recommends that you use the Configuration utility to create and manage customization groups.

SEE ALSO

apm policy agent, apm profile

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy customization-group(1)

apm policy customization-languages

NAME

customization-languages - Deprecated singleton container for customization languages

MODULE

apm policy

DESCRIPTION

customization-languages has been deprecated. It has been replaced by the accept-languages keyword in apm profile access customization profiles.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2017. All rights reserved.

BIG-IP 2017-09-05 apm policy customization-languages(1)

apm policy image-file

NAME

image-file - Manages a file that contains an image.

MODULE

apm policy

SYNTAX

Warning: F5 Networks recommends that you use the Configuration utility to create and manage image files.

SEE ALSO

apm policy agent, apm profile

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-19 apm policy image-file(1)

apm policy policy-item

NAME

policy-item - Manages an access policy item.

MODULE

apm policy

SYNTAX

Warning: F5 Networks recommends that you use the visual policy editor in the Configuration utility to create and manage access policy items.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-03-22 apm policy policy-item(1)

apm policy windows-group-policy-file

NAME

windows-group-policy-file - Manages FullArmor GPAnywhere Windows group policy files.

MODULE

apm policy

SYNTAX

Warning: This page is obsolete. Windows Group Policy is no longer supported.

Warning: F5 Networks recommends that you use the visual policy editor in the Configuration utility to create and manage FullArmor GPAnywhere Windows group policy files.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2016-03-14 apm policy windows-group-policy-file(1)

apm profile access

NAME

access - Configures an access profile.

MODULE

apm profile

SYNTAX

Configure the access component within the profile module using the syntax shown in the following sections.

CREATE/MODIFY

create access [name]

options:

```
accept-languages [add | delete | modify | replace-all-with] {
  [name]
}
access-policy [[string] | none]
access-policy-timeout [integer]
app-service [[string] | none]
cache-generation [integer]
customization-group [[string] | none]
default-language [[string] | none]
defaults-from [[string] | none]
domain-cookie [[string] | none]
domain-groups [add | delete | modify | replace-all-with] {
  [name]
}
domain-mode [single-domain | multi-domain]
user-identity-method [http | ip-address]
enforce-policy [true | false]
eps-group [[string] | none]
errormap-group [[string] | none]
framework-installation-group [[string] | none]
general-ui-group [[string] | none]
generation-action [increment | noop]
httponly-cookie [true | false]
inactivity-timeout [integer]
logout-uri-include [add | delete | modify | replace-all-with] {
  [name]
}
logout-uri-timeout [integer]
log-settings [add | delete | modify | replace-all-with] {
  [name]
}
max-concurrent-sessions [[integer] | none]
max-concurrent-users [[integer] | none]
max-failure-delay [integer]
max-in-progress-sessions [[integer] | none]
max-session-timeout [integer]
min-failure-delay [integer]
oauth-profile [[oauth-profile-name] | none]
persistent-cookie [true | false]
primary-auth-service [[string] | none]
restrict-to-single-client-ip [true | false]
sandboxes [add | delete | modify | replace-all-with] {
  [name] { retain-public-access [true|false] }
}
scope [profile | virtual-server | global | named | public]
named-scope [[string] | none]
secure-cookie [true | false]
sso-name [[string] | none]
type [all | identity-service | ltm-apm | oauth-resource-server | rdg-rap | ssl-vpn | sso | swg-explicit | swg-transparent | system-authenticatio
use-http-503-on-error [true | false]
```

modify access [name]

options:

```

accept-languages [add | delete | modify | replace-all-with] {
  [name]
}
access-policy [[string] | none]
access-policy-timeout [integer]
app-service [[string] | none]
cache-generation [integer]
customization-group [[string] | none]
default-language [[string] | none]
defaults-from [[string] | none]
domain-cookie [[string] | none]
domain-groups [add | delete | modify | replace-all-with] {
  [name]
}
domain-mode [single-domain | multi-domain]
user-identity-method [http | ip-address]
enforce-policy [true | false]
eps-group [[string] | none]
errormap-group [[string] | none]
framework-installation-group [[string] | none]
general-ui-group [[string] | none]
generation-action [increment | noop]
httponly-cookie [true | false]
inactivity-timeout [integer]
logout-uri-include [add | delete | modify | replace-all-with] {
  [name]
}
logout-uri-timeout [integer]
log-settings [add | delete | modify | replace-all-with] {
  [name]
}
max-concurrent-sessions [[integer] | none]
max-concurrent-users [[integer] | none]
max-failure-delay [integer]
max-in-progress-sessions [[integer] | none]
max-session-timeout [integer]
min-failure-delay [integer]
oauth-profile [[oauth-profile-name] | none]
persistent-cookie [true | false]
primary-auth-service [[string] | none]
restrict-to-single-client-ip [true | false]
sandboxes [add | delete | modify | replace-all-with] {
  [name] { retain-public-access [true|false] }
}
scope [profile | virtual-server | global | named | public]
named-scope [[string] | none]
secure-cookie [true | false]
sso-name [[string] | none]
use-http-503-on-error [true | false]
edit access [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

DISPLAY

```

list access
list access [ [ [name] | [glob] | [regex] ] ... ]
show running-config access
show running-config access [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition

```

```

show access
show access [name]

```

DELETE

```

delete access [name]

```

DESCRIPTION

You can use the access component to configure an access profile. An access profile is a pre-configured group of settings that you can use to configure secure Network Access for an application.

EXAMPLES

```

create access MyAccessProfile { defaults-from access access-policy "my_access_policy" accepted-languages
  "my_accepted_languages" default-language "en" customization-group "company_logout" eps-group 'myepsgroup'
  framework-installation-group "company_header" "company_footer" errormap-group "company_errormap" }

```

Creates an access profile named MyAccessProfile that is based on the default access profile named access, uses the access policy named my_access_policy, accepts the languages in the my_accepted_languages class, uses English as the default language, and uses these groups to customize the application pages and messages: company_logout, company_header, company_footer, and company_errormap.

```

list access all all-properties
Displays a list of access profiles, including parameter values.

```

```

delete access MyAccessProfile

```

Deletes the access profile named MyAccessProfile.

OPTIONS

accept-languages

Specifies the name of a class that defines the languages supported by the access profile. The default languages are en (English), ja (Japanese), zh-cn (simplified Chinese (PRC)), and zh-tw (traditional Chinese (Taiwan)). This option is required.

access-policy

Specifies the access policy that you want to enforce using this access profile. An access policy contains various security checks that a client must pass before the BIG-IP Access Policy Manager grants access to a protected application. This option is required.

access-policy-timeout

Specifies, for this access profile, the number of seconds within which a user must complete the steps to gain access to an application. The default is 300 seconds. This option is designed to quickly release session resources when a user does not complete the access process, for example, when the user closes the browser before completing the access process.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

customization-group

Specifies the customization group that defines the appearance of the logout and error pages. This option is required.

default-language

Specifies the default language for the BIG-IP Access Policy Manager that you want to implement with this access profile. The default is en (English). If the client requests a language that is not supported, the BIG-IP Access Policy Manager uses the default value. This option is required.

defaults-from

Specifies the default access policy from which this profile is created. This option is required.

domain-cookie

Specifies a domain cookie to use with an application access control connection. If you specify a domain cookie, then the line domain=specified_domain is added to the MRHsession cookie. The default is none.

domain-groups

Specifies a group of multiple domains or multiple hosts in multiple domains to which a single user session has access. For example, you can use this option to configure a single user session to have access to three domains: www.a.com, www.b.com, and www.c.com. When a user logs in to any of these domains, that user can access the other domains without logging in again. This option is required when you set the domain-mode option to multi-domain. This option is ignored when you set the access domain-mode option to single-domain.

For each domain in the domain group, you can specify the following settings:

cookie-host

Specifies the host name for which to create the user's session cookie.

cookie-domain

Specifies the domain for which to create the user's session cookie.

secure-cookie

Adds a security attribute to the user's session cookie.

persistent-cookie

Adds a persistence attribute to the user's session.

sso-name

Specifies the SSO method to use when accessing a backend application.

domain-mode

Specifies how the SSO configuration is applied. The options are:

single-domain

Applies the SSO configuration to a single domain. This is the default.

When you set domain-mode to single-domain, you must also set the sso-name option.

multi-domain

Applies the SSO configuration across multiple domains. This option allows users a single APM login/session and applies the credentials across multiple Local Traffic Manager or Access Policy Manager virtual servers in front of different domains. Note that to apply SSO configurations across multiple domains, all virtual servers must be on one BIG-IP system.

When you set domain-mode to multi-domain, you must also configure the domain-group option, and provide a URI for the primary-auth-service option.

user-identity-method

Specifies how access will bind a session to a request.

http Use http information such as cookies and URI query string to identify user.

ip-address

Use IP address to identify a user. Do not use this setting if clients may be behind a NAT.

enforce-policy

Set this option to false, if you don't want to enforce the access-policy. The default is true which means the access-policy is always enforced. This option can only be modified for SWG-Transparent type profile.

eps-group

This option is required.

errormap-group

Specifies the customization settings for the error map that you want to implement with this access profile. This setting is required.

framework-installation-group

Specifies the customization settings for the header and footer that you want to implement with this access profile. This setting is required.

generation-ui-group

Specifies the generation of the user interface group for the new generation access configuration. This option is required.

generation-timeout

Specifies the timeout, in seconds, for the new generation access configuration.

generation-action

increment

Activates the current access policy configuration for an access profile. For example, the following command activates current access policy configuration for profile myAccessProfile: `tms modify apm profile access myAccessProfile generation-action increment`

`noop` Specifies "no operation to be performed". This is the default.

`sync` Specifies that the policy is being modified due to APM policy sync operation. This is an internal action; you should not set it.

httponly-cookie

Specifies whether HttpOnly directive should be inserted in HTTP response from BIG-IP. The client browser should prevent script from accessing cookie, if this flag is set in the response. The default is false.

inactivity-timeout

Specifies, for this access profile, the number of seconds that the session on the client can be idle before the server disconnects the VPN tunnel. The default is 900 seconds.

logout-uri-include

Specifies a list of URIs to include in the access profile for initiating session logout.

logout-uri-timeout

Specifies the timeout used to delay logout for the customized logout URIs defined in the logout uri include list

log-settings

Specifies one or more log-setting containers to associate with this profile

max-concurrent-sessions

Specifies, for this access profile, the number of concurrent sessions allowed. The default is 0 (zero), which represents unlimited sessions. Users assigned an administrative role of Application Editor can view the value of this option. Users assigned any other administrative role can modify this option.

max-concurrent-users

Specifies, for this access profile, the number of concurrent sessions allowed. The default is 0 (zero), which represents unlimited sessions. This field is Read-only for Application Editors. Users assigned any other administrative role can modify this field.

max-failure-delay

Specifies the maximum random delay after authentication failure during the access policy. It is the maximum number of seconds before the user is shown an error message on the logon page and prompted to re-enter credentials. The default is 5 seconds. 0 (zero) represents no delay. Note: Set max-failure-delay to no more than one-half the access-policy-timeout value and no more than 65 seconds greater than min-failure-delay.

max-in-progress-sessions

Specifies the maximum number of in-progress concurrent sessions a user can have. The in-progress sessions are the sessions for which an access policy has not completed. The default is 0, which represents an unlimited number of such sessions.

max-session-timeout

Specifies the maximum lifetime of one session. The maximum lifetime is the number of seconds between session creation and session termination.

min-failure-delay

Specifies the minimum random delay after authentication failure during the access policy. It is the minimum number of seconds before the user is prompted for credentials again or shown an error message on the logon page. The default is 2 seconds.

[name]

Specifies the name of the access profile. This option is required.

oauth-profile

Specifies an oauth profile for use with an OAuth Authorization Server.

persistent-cookie

Specifies to retain the cookie for a user session, even when the user session is terminated, when set to true. Although this is an insecure method, this setting is useful and required in cases where you have a third-party application, such as Sharepoint, and need to store the cookie in a local database so that any attempt to access backend server applications through Access Policy Manager succeeds. The default is false.

primary-auth-service

Specifies the address of your primary authentication URI. This setting is required when you set the domain-mode option to multi-domain.

For example, when you set this option to `https://logon.yourcompany.com`, the user session is stored on this primary domain, and the user can access multiple backend applications from multiple domains and hosts without re-entering credentials.

restrict-to-single-client-ip

Specifies whether a user session is tied to a single client IP. If during session's lifetime, the user's client IP address changes, the current session is terminated. The user needs to re-login to create a new session from the new client IP address. The default is false.

sandboxes

Specifies the association between the access profile and the sandbox. If `retain-public-access` is set to true, this association is retained even if there is no resource that uses sandbox files in the access policy that corresponds to this access profile.

scope

Specifies the confining scope for sessions created by the profile. Set this option to profile (which is also the default-value) to confine the validity of a session to the profile from which it was created. Set this option to virtual-server to further confine the validity of a session to the virtual server from which it was created. Setting this option to global allows the session to be valid on any virtual server with any access profile that also specifies global scope. Setting this option to named allows the session to be valid for any virtual server with access profile using the same named-scope value. The option public is allowed for only SSLO access profiles and sessions aren't created.

named-scope

Specifies the string to which the validity of a session is confined to. This setting is required when you set the scope option to named.

secure-cookie

Set this option to true, if you want to add a secure keyword to the session cookie. Set this option to false, if you want to configure an application access control scenario that uses an HTTPS virtual server to authenticate the user, and then sends the user to an existing HTTP virtual server to use applications. The default is true.

sso-name

Specifies the SSO configuration that you want BIG-IP Access Policy Manager to use to submit the user's credentials to the backend application. This allows the user to log in once to the Access Policy Manager and then gain access to backend applications without logging in again.

`type` Specifies the type of access profile. You can specify the following types for an access profile.

`all` Supports ltm-apm and ssl-vpn access types.

identity-service

Used internally to provide identity service for a supported integration. Only APM creates this type of profile.

ltm-apm

For web access management configuration.

oauth-resource-server

Supports apps and devices that use OAuth tokens but do not support cookies.

rdg-rap

For validating connections to hosts behind APM when APM acts as a gateway for RDP clients.

ssl-vpn

For network access, portal access, or application access.

`sso` For configuring matching virtual servers for Single Sign-On (SSO).

swg-explicit

For Secure Web Gateway explicit forward proxy.

swg-transparent

For Secure Web Gateway transparent forward proxy.

system-authentication

For configuring administrator access to the BIG-IP system (when using APM as a pluggable authentication module).

use-http-503-on-error

Set this option to true to use HTTP response code 503 for error pages sent by BIG-IP Access Policy

Manager to clients. Set this option to false to use HTTP response code 200. The default is false.

SEE ALSO

apm sso, apm policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2019-02-10 apm profile access(1)

apm profile connectivity

NAME

connectivity - Configures a connectivity profile.

MODULE

apm profile

SYNTAX

Configure the connectivity component within the profile module using the syntax shown in the following sections.

CREATE/MODIFY

create connectivity [name]

modify connectivity [name]

options:

adaptive-compression [enabled | disabled]

app-service [[string] | none]

citrix-client-bundle [[name] | default-citrix-client-bundle]

client-policy [add | delete | modify | replace-all-with] {
[name] {

android-ec {

device-lock-method [alphabetic | alphanumeric | any | numeric]

enable-mobilesafe [true | false]

enforce-device-lock [true | false]

enforce-logon-mode [true | false]

logon-mode [native | web]

require-device-auth [true | false]

max-inactivity-time [integer]

min-passcode-length [integer]

save-password [true | false]

save-password-method [disk | memory]

save-password-timeout [integer]

}

android-ep {

device-lock-method [alphabetic | alphanumeric | any | numeric]

enable-mobilesafe [true | false]

enforce-device-lock [true | false]

enforce-logon-mode [true | false]

logon-mode [native | web]

max-inactivity-time [integer]

min-passcode-length [integer]

save-password [true | false]

save-password-method [disk | memory]

save-password-timeout [integer]

}

chromeos-ec {

enforce-logon-mode [true | false]

logon-mode [native | web]

save-password [true | false]

save-password-method [disk | memory]

save-password-timeout [integer]

}

macos-ec {

enforce-logon-mode [true | false]

logon-mode [native | web]

save-password [true | false]

save-password-method [disk | memory]

save-password-timeout [integer]

}

ec {

component-update [yes | prompt | no]

location-dns [add | delete | modify | replace-all-with] {

[name]

}

```

reuse-winlogon-creds [true | false]
reuse-winlogon-session [true | false]
save-password [true | false]
save-password-method [disk | memory]
save-password-timeout [integer]
save-servers-on-exit [true | false]
}
ios-ec {
enable-mobilesafe [true | false]
enforce-logon-mode [true | false]
logon-mode [native | web]
require-device-auth [true | false]
save-password [true | false]
save-password-method [disk | memory]
save-password-timeout [integer]
vod-disconnect-timeout [integer]
}
ios-ep {
enable-mobilesafe [true | false]
enforce-logon-mode [true | false]
logon-mode [native | web]
enforce-pin-lock [true | false]
max-grace-period [integer]
save-password [true | false]
save-password-method [disk | memory]
save-password-timeout [integer]
}
oauth {
provider-name [name]
client-id [string]
scopes [string]
done-uri [string]
}
servers {
{
alias [[string] | none]
host [string]
}
...
}
}
compress-buffer-size [integer]
compress-cpu-saver [true | false]
compress-cpu-saver-high [integer]
compress-cpu-saver-low [integer]
compress-gzip-level [integer]
compress-gzip-memlevel [integer]
compress-gzip-window-size [integer]
compress-ingress [true | false]
compress-preferred-method [[string] | none]
compression [enabled | disabled]
compression-codecs [[string] | none]
customization-group [[string] | none]
defaults from [[name] | none]
deflate-compression-level [integer]
description [[string] | none]
location-specific [true | false]
tunnel-name [[string] | none]

```

edit connectivity [[[name] | [glob] | [regex]] ...]

options:

```

all-properties
non-default-properties

```

DISPLAY

list connectivity

list connectivity [[[name] | [glob] | [regex]] ...]

show running-config connectivity

show running-config connectivity [[[name] | [glob] | [regex]] ...]

options:

```

all-properties
non-default-properties
partition

```

show connectivity

show connectivity [name]

DELETE

delete connectivity [name]

DESCRIPTION

You can use the connectivity component to configure a connectivity profile. By using the connectivity profile, you can configure L2 and L4 tunnels, compression, Windows and mobile client settings, and client component downloads from F5 Networks and Citrix.

EXAMPLES

create connectivity myconnectivityprofile { }

Creates a connectivity profile named myconnectivityprofile that inherits its settings from the system default connectivity profile.

OPTIONS

adaptive-compression

Enables or disables adaptive compression. Use this option to configure compression settings for application tunnels and to optimize applications and RDP traffic. The default is enabled.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

citrix-client-bundle

Specifies the Citrix client bundle used by this connectivity profile. The default is default-citrix-client-bundle.

client-policy

Adds, deletes, or modifies the client policy for any of the following clients:

android-ec Android Edge Client

android-ep Android Edge Portal

chromeos-ec Chrome OS Edge Client

macos-ec Mac OS F5 Access

ec Windows/OSX Edge Client

ios-ec iOS Edge Client

ios-ep iOS Edge Portal

Options (please refer to the SYNTAX section to see if a certain option is supported for a particular client):

component-update

Specifies how the client handles automatic updates. The options are:

yes Automatically installs a client update when one is available.

prompt

Prompts the user about installing a client update.

No Disables the client from receiving automatic updates.

device-lock-method

Specifies the device lock quality that the client should enforce on the device. The options are:

alphabetic

Device passcode must contain at least alphabetic (or other symbol) characters.

alphanumeric

Device passcode must contain at least both numeric and alphabetic (or other symbol) characters.

any A device passcode must be set but does not matter what it is.

numeric

Device passcode must contain at least numeric characters.

enable-mobilesafe

Enables or disable MobileSafe checks. Use this option to configure whether client should execute the MobileSafe security checks as part of the logon. The default is false.

enforce-device-lock

Specifies whether client should enforce a device passcode policy on the device. The default is true.

enforce-logon-mode

Specifies whether client should enforce a logon mode on the device. The default is false. Set to true if external logon page is used.

logon-mode

Specifies logon mode to be enforced on the device. The default is native. Set to web if external logon page is used.

enforce-pin-lock

Specifies whether client should enforce an app-level PIN before allowing access to the app. The default is true.

location-dns

Specifies a list of DNS suffixes used by the Network Location Awareness feature of the client. This list represents the internal network where local resources are available without the need of a Network Access connection. The default is none.

`max-grace-period`
Specifies the length of time (in minutes) the app was taken to the background before the user will be asked for a PIN. With the option set to 0, user will be asked for the PIN every time the app is taken from the background. The default is 2.

`max-inactivity-time`
Sets the length of time (in minutes) since the user last touched the screen or pressed a button before the device locks the screen. The default is 5.

`min-passcode-length`
Specifies the minimum required number of characters for the device passcode. The default is 4.

`oauth`
OAuth configuration for EDGE clients.

`client-id`
Specifies OAuth client identifier. The client identifier is not a secret; it is exposed by BIG-IP APM virtual server. OAuth configuration is disabled if `client-id` is not specified. The default is none.

`done-uri`
Specifies URI for OAuth client to be directed to when authentication complete or failed ("You can close this tab" page). Default APM page is used when none is selected. The default is none.

`provider-name`
Specifies the name of the OAuth provider (`apm aaa oauth-provider`). OAuth configuration is disabled if none is not specified. The default is none.

`scopes`
Specifies scope of the OAuth access request. The value of the `scopes` parameter is expressed as a list of space-delimited, case-sensitive strings. The strings are defined by the authorization server. If the value contains multiple space-delimited strings, their order does not matter. Only 0x20-0x21, 0x23-0x5B, 0x5D-0x7E characters are allowed. The default is none.

`require-device-auth`
Specifies whether device authentication is needed before accessing cached credentials. The default is false.

`reuse-winlogon-creds`
Specifies whether client can reuse logon credentials entered by a user for a subsequent log in. The default is false.

`reuse-winlogon-session`
Specifies whether client should attempt to use the same Windows logon session. The default is false.

`save-password`
Specifies whether client allows user password caching. The default is false.

`save-password-method`
Specifies whether client saves encrypted passwords on disk or caches passwords in memory only. The default is disk.

`save-password-timeout`
Specifies the number of minutes that a cached password remains valid (applies only to in-memory password caching). The default is 240.

`save-servers-on-exit`
Specifies whether client maintains a list of Access Policy Manager systems that the client accessed. The default is true.

`servers`
Specifies a list of server and alias pairs in the client's server list.

`compress-buffer-size`
Specifies the size of compressed data for Network Access tunnels. The default is 4096.

`compress-cpu-saver`
Specifies whether the system monitors the percentage of CPU usage and adjusts compression rates automatically when CPU usage reaches either the CPU saver high threshold or the CPU saver low threshold. The default is true.

`compress-cpu-saver-high`
Specifies the percentage of CPU usage at which the system starts automatically decreasing the amount of content being compressed, as well as the amount of compression which the system is applying. The default is 90 percent.

`compress-cpu-saver-low`
Specifies the percentage of CPU usage at which the system resumes content compression at the user-defined rates. The default is 75 percent.

`compress-gzip-level`
Specifies the degree to which the system compresses the content. Higher compression levels slow down the compression process. The default is 6, which provides a higher amount of compression at the expense of more CPU processing time. 1 is the lowest level and 9 is the highest level. 0 disables compression.

`compress-gzip-memlevel`

Specifies the number of kilobytes of memory that the system uses for internal compression buffers when compressing data. You can select a value between 1 and 256. The default is 8192.

`compress-gzip-window-size`

Specifies the number of kilobytes in the window size that the system uses when compressing data. You can select a value between 1 and 128. The default is 16384.

`compress-ingress`

Specifies whether incoming data is compressed. The default is false.

`compress-preferred-method`

Specifies the preferred method of data compression. The default is zlib.

`compression`

Enables or disables compression between the client and the server. The default is enabled.

`compression-codecs`

Specifies the available compression codecs for server-to-client connections. The server compares the available compression types you configure with the available compression types on the client, and then chooses the most effective mutual compression setting. Compression for the client is configured separately. The default includes all three available codecs:

`lzo` Offers a balance between CPU resources and compression ratio, compressing more than deflate, but with less CPU resources than `bzip2`.

`deflate`

Uses the least CPU resources, but compresses the least effectively.

`bzip2`

Uses the most CPU resources, but compresses the most effectively.

`customization-group`

Specifies which customization groups are applied. This option is required.

`defaults-from`

Specifies the profile from which this profile inherits properties that are not specified explicitly. The default is connectivity.

`deflate-compression-level`

Specifies the level of compression performed by the deflate codec. The default is 1.

`description`

Specifies a user-defined description for the connectivity profile.

`location-specific`

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

`[name]`

Specifies the profile that you want to use as the parent profile. The new profile inherits all settings and values from the parent profile.

`tunnel-name`

Specifies the name of the tunnel through which data passes. The default is none.

SEE ALSO

`apm aaa oauth-provider`, `apm profile`, `ltm virtual`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2020-01-31 `apm profile connectivity(1)`

apm profile exchange

NAME

`exchange` - Configures an exchange profile.

MODULE

`apm profile`

SYNTAX

Configure the exchange component within the profile module using the syntax shown in the following sections.

CREATE/MODIFY

create exchange [name]

modify exchange [name]

options:

ntlm-auth-name [[string] | none]
active-sync-url [[string] | none]
active-sync-auth-type [basic | ntlm | basic-ntlm]
active-sync-sso-config [[string] | none]
auto-discover-url [[string] | none]
auto-discover-auth-type [basic | ntlm | basic-ntlm]
auto-discover-sso-config [[string] | none]
description [[string] | none]
offline-address-book-url [[string] | none]
offline-address-book-auth-type [basic | ntlm | basic-ntlm]
offline-address-book-sso-config [[string] | none]
rpc-over-http-url [[string] | none]
rpc-over-http-auth-type [basic | ntlm | basic-ntlm]
rpc-over-http-sso-config [[string] | none]
user-agent-pattern-for-utf8 [[string] | none]
web-service-url [[string] | none]
web-service-auth-type [basic | ntlm | basic-ntlm]
web-service-sso-config [[string] | none]

edit exchange [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list exchange

list exchange [[[name] | [glob] | [regex]] ...]

show running-config exchange

show running-config exchange [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
partition

show exchange

show exchange [name]

DELETE

delete exchange [name]

DESCRIPTION

You can use the exchange component to configure an exchange profile. An exchange profile is a preconfigured group of settings that you can use to configure authentication for exchange services such as Outlook Anywhere, ActiveSync, Autodiscover and Offline Address Book, so that those work with BIG-IP.

EXAMPLES

```
create exchange MyExchangeProfile {  
  ntlm-auth-name "MyNTLMAuth"  
  rpc-over-http-url "/rpc/rpcproxy.dll"  
  rpc-over-http-auth-type ntlm  
  rpc-over-http-sso-config "MyKerberosSSOConfig"  
}>
```

Creates an exchange profile named MyExchangeProfile that is based on the general settings such as NTLM Authentication configuration MyNTLMAuth. The profile is configured for Outlook Anywhere (RPC over HTTP) service with url "/rpc/rpcproxy.dll", client authentication type ntlm and SSO configuration type MyKerberosSSOConfig

```
list exchange all all-properties
```

Displays a list of exchange profiles, including parameter values.

```
delete access MyExchangeProfile
```

Deletes the exchange profile named MyExchangeProfile.

OPTIONS

ntlm-auth-name

Specifies the NTLM configuration object that clients can use to authenticate on the front-end. Backend SSO type must be Kerberos for ntlm or basic-ntlm front end.

active-sync-auth-type

Specifies the client-side authentication type for ActiveSync exchange service. The valid value is basic.

active-sync-sso-config

Specifies the back end SSO config for ActiveSync exchange service. This is optional.

active-sync-url

Specifies the URL for ActiveSync exchange service. URL is required for ActiveSync service to be enabled through BIG-IP.

auto-discover-auth-type

Specifies the client-side authentication type for Autodiscover exchange service. The valid values are basic, ntlm and basic-ntlm.

`auto-discover-sso-config`
Specifies the back end SSO config for Autodiscover exchange service. This is optional.

`description`
Specifies a user-defined description for the exchange profile.

`auto-discover-url`
Specifies the URL for Autodiscover exchange service. URL is required for Autodiscover service to be enabled through BIG-IP.

`offline-address-book-auth-type`
Specifies the client-side authentication type for Offline Address Book exchange service. The valid values are basic, ntlm and basic-ntlm.

`offline-address-book-sso-config`
Specifies the back end SSO config for Offline Address Book exchange service. This is optional.

`offline-address-book-url`
Specifies the URL for Offline Address Book exchange service. URL is required for Offline Address Book service to be enabled through BIG-IP.

`rpc-over-http-auth-type`
Specifies the client-side authentication type for Outlook Anywhere (RPC over HTTP) exchange service. The valid values are basic, ntlm and basic-ntlm.

`rpc-over-http-sso-config`
Specifies the back end SSO config for Outlook Anywhere (RPC over HTTP) exchange service. This is optional.

`rpc-over-http-url`
Specifies the URL for Outlook Anywhere (RPC over HTTP) exchange service. URL is required for Outlook Anywhere (RPC over HTTP) service to be enabled through BIG-IP.

`user-agent-pattern-for-utf8`
Specifies the user agent pattern for UTF8.

`web-service-auth-type`
Specifies the client-side authentication type for Web Exchange service. The valid values are basic, ntlm and basic-ntlm.

`web-service-sso-config`
Specifies the back end SSO config for Web Exchange service. This is optional.

`web-service-sync-url`
Specifies the URL for Web Exchange service. URL is required for Web Service to be enabled through BIG-IP.

SEE ALSO

`apm sso`, `apm profile access`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2015-05-12 `apm profile exchange(1)`

apm profile oauth

NAME
`oauth` - Configures an oauth profile.

MODULE
`apm profile`

SYNTAX
Configure the oauth component within the profile module using the syntax shown in the following sections.

CREATE/MODIFY
`create oauth [name]`
`modify oauth [name]`
options:
`access-token-lifetime [integer]`
`app-service [[string] | none]`
`audience [add | delete | none | replace-all-with] {`
`[string]`
`}`

```

auth-code-lifetime [integer]
auth-url [string]
client-apps [add | delete | replace-all-with] {
  [client-app-name]
}
db-instance [db-instance-name]
defaults-from [[string] | none]
generate-jwt-refresh-token [true | false]
generate-refresh-token [true | false]
id-token-claims [add | delete | none | replace-all-with] {
  [claim-name]
}
id-token-lifetime [integer]
id-token-primary-key [jwk-config-name]
ignore-expired-cert [true | false]
issuer [string]
jwks-url [string]
jwt-access-token-claims [add | delete | none | replace-all-with] {
  [claim-name]
}
jwt-access-token-lifetime [integer]
jwt-ec-signature-format [binary | der]
jwt-refresh-token-enc-secret [string]
jwt-refresh-token-lifetime [integer]
jwt-token [enabled | disabled]
opaque-token [enabled | disabled]
openid-cfg-url [string]
openid-connect [enabled | disabled]
per-user-token-limit [integer]
primary-key [jwk-config-name]
refresh-token-lifetime [integer]
refresh-token-usage-limit [integer]
resource-servers [add | delete | replace-all-with] {
  [resource-server-name]
}
reuse-access-token [true | false]
reuse-refresh-token [true | false]
rotation-keys [add | delete | none | replace-all-with] {
  [jwk-config-name]
}
subject [[string] | none]
token-introspection-url [string]
token-issuance-url [string]
token-revocation-url [string]
trusted-ca-bundle [certificate-file-object-name]
userinfo-claims [add | delete | none | replace-all-with] {
  [claim-name]
}
userinfo-primary-key [jwk-config-name]
userinfo-url [string]
edit oauth [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

DISPLAY

```

list oauth
list oauth [ [ [name] | [glob] | [regex] ] ... ]
show running-config oauth
show running-config oauth [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition

```

```

show oauth
show oauth [name]

```

DELETE

```

delete oauth [name]

```

DESCRIPTION

You can use the oauth component to configure an oauth profile. An oauth profile is a pre-configured group of settings that you can use to configure OAuth Authorization Server.

NOTE: For the oauth profile to take effect, this profile must be associated with an access profile. (See man page for apm access profile.)

EXAMPLES

```

create oauth myOAuthProfile {
  defaults-from oauth
  client-apps add { client_1 client_2 }
  resource-servers add { rs_1 rs_2 }
  opaque-token enabled
  db-instance db_test
  jwt-token enabled
  openid-connect enabled
}

```

```

issuer https://example.f5.com
primary-key jwk1_hs256
id-token-primary-key jwk1_rs256
generate-jwt-refresh-token true
jwt-refresh-token-enc-secret password
auth-url /f5-oauth2/v1/authorize
token-issuance-url /f5-oauth2/v1/token
token-revocation-url /f5-oauth2/v1/revoke
token-introspection-url /f5-oauth2/v1/introspect
openid-cfg-url /f5-oauth2/v1/.well-known/openid-configuration
jwks-url /f5-oauth2/v1/jwks
userinfo-url /f5-oauth2/v1/userinfo
}

```

Creates an oauth profile named myOAuthProfile that is based on the default oauth profile named oauth. The profile serves OAuth requests from client applications named client_1 and client_2 and resource servers named rs_1 and rs_2.

The profile is configured to generate both Opaque and JWT access tokens. For Opaque access token, it uses db instance named db_test. For JWT access token, it uses issuer named https://example.f5.com, primary key named jwk1_hs256 to sign JWT tokens and JWT refresh token encryption secret named password for encryption of refresh token generated with the JWT access token. The profile also supports OpenID Connect. It uses key named jwk1_rs256 to sign ID Tokens.

It uses /f5_oauth2/v1/authorize as the authorization endpoint, /f5-oauth2/v1/token as token issuance endpoint, /f5-oauth2/v1/revoke as revocation endpoint, /f5-oauth2/v1/introspect as token introspection endpoint for validating Opaque tokens, /f5-oauth2/v1/.well-known/openid-configuration as OpenID Connect metadata configuration endpoint, /f5-oauth2/v1/jwks as JWKS endpoint and /f5-oauth2/v1/userinfo as UserInfo endpoint.

list oauth all all-properties

Displays a list of oauth profiles, including parameter values.

delete oauth myOAuthProfile

Deletes the oauth profile named myOAuthProfile.

OPTIONS

access-token-lifetime
Specifies the number of minutes for which the access token should be valid. The default is 5 minutes.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

audience
Specifies the audience claim for which the JWT access token is intended. This is a list of values. Each value in this list can be a string, URI, or session variable.

auth-code-lifetime
Specifies the number of minutes for which the authorization code should be valid. The default is 5 minutes.

auth-url
Specifies the path of the authorization endpoint that is used to authenticate the resource owner and provide the authorization code. The default is /f5-oauth2/v1/authorize.

client-apps
Specifies the list of client applications that is served by the OAuth Authorization Server associated with this profile.

db-instance
Specifies the db instance that is used to store tokens generated by the OAuth Authorization Server that is associated with this profile.

defaults-from
Specifies the default oauth profile from which this profile is created. The default is oauth.

generate-jwt-refresh-token
Specifies whether a refresh token should be generated along with the JWT access token. This is applicable only for "Authorization Code" and "Resource Owner Password Credentials" grant types. The default is true.

generate-refresh-token
Specifies whether a refresh token should be generated along with the access token. This is applicable only for "Authorization Code" and "Resource Owner Password Credentials" grant types.

id-token-claims
Specifies the list of claims that are part of ID token.

id-token-lifetime
Specifies the number of minutes for which the ID token should be valid. The default is 5 minutes.

id-token-primary-key
Specifies the JWK config that is used to retrieve the shared key (symmetric) or private key (asymmetric) used to sign ID token. If the key is asymmetric, the configured public key will be returned as part of

JWKS URL response.

`ignore-expired-cert`

Specifies whether to ignore the expiry of the certificate used for signing JWT access token. If this value is true, then the certificate will be used for signing JWT access token even if it is expired. The default is false.

`issuer`

Specifies the issuer claim that is part of JWT access token. This value must be a URI.

`jwt-key`

Specifies the path of the JWKS endpoint that returns public signing keys. These keys are used by OAuth Resource Servers to verify the digital signature of JWT access token. The default is `/f5-oauth2/v1/jwks`.

`jwt-access-token-claims`

Specifies the list of claims that are part of JWT access token.

`jwt-access-token-lifetime`

Specifies the number of minutes for which the JWT access token should be valid. The default is 5 minutes.

`jwt-ec-signature-format`

Specifies the JWT token signature format for Elliptic Curve. The default is binary format.

`jwt-refresh-token-enc-secret`

Specifies the JWT refresh token encryption secret that is used to generate an encryption key. This key is used to encrypt the refresh token when JWT token is enabled.

`jwt-refresh-token-lifetime`

Specifies the number of minutes for which the JWT refresh token should be valid. The default is 60 minutes.

`jwt-token`

Specifies whether JWT access token should be generated. The default is false.

`opaque-token`

Specifies whether opaque (non-JWT) access token should be generated. The default is true.

`openid-cfg-url`

Specifies the path of OpenID Connect endpoint that returns OpenID Connect configuration. The default is `/f5-oauth2/v1/.well-known/openid-configuration`.

`openid-connect`

Specifies whether this OAuth profile supports OpenID connect or not.

`per-user-token-limit`

Specifies the maximum number of active access tokens that can be generated for a user. The default is 255. The range is 0 to 5000.

`primary-key`

Specifies the JWK config that is used to retrieve the shared key (symmetric) or private key (asymmetric) used to sign JWT access token. If the key is asymmetric, the configured public key will be returned as part of JWKS URL response.

`refresh-token-lifetime`

Specifies the number of minutes for which the refresh token should be valid. The default is 480 minutes.

`refresh-token-usage-limit`

Specifies the maximum number of times the access token can be obtained using the refresh token request. The default value is 0, which represents unlimited number of times.

`resource-servers`

Specifies the list of resource servers that is served by the OAuth Authorization Server that is associated with this profile.

`reuse-access-token`

Specifies whether an access token is reused or a new access token is generated when it is obtained using refresh token request. When the access token is reused, its expiry time is extended.

`reuse-refresh-token`

Specifies whether a refresh token is reused or a new refresh token is generated when it is obtained using refresh token request.

`rotation-keys`

Specifies one or more JWK configs that contain public keys used as rotation keys. The public keys derived from this set will be returned as part of JWKS URL response.

`subject`

Specifies the subject claim that is part of JWT access token. This value can be a string, URI, or session variable. The default is `#{session.assigned.uid}`

`token-issuance-url`

Specifies the path of token issuance endpoint that is used to issue an access token and possibly a refresh token. The default is `/f5-oauth2/v1/token`.

`token-revocation-url`

Specifies the path of token revocation endpoint that is used to revoke an access token or a refresh token. The default is `/f5-oauth2/v1/revoke`.

token-introspection-url

Specifies the path of token introspection endpoint that is used to introspect an access token. The default is /f5-oauth2/v1/introspect.

trusted-ca-bundle

Specifies the trusted ca bundle that is used during verification of JWK config specified in primary-key that uses asymmetric key.

userinfo-claims

Specifies the list of claims that are part of UserInfo.

userinfo-primary-key

Specifies the JWK config that is used to retrieve the shared key (symmetric) or private key (asymmetric) used to sign UserInfo. If the key is asymmetric, the configured public key will be returned as part of JWKS URL response.

userinfo-url

Specifies the path of userinfo endpoint that is used to obtain claims about the authenticated end-user. The default is /f5-oauth2/v1/userinfo.

SEE ALSO

apm oauth, apm policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015-2017. All rights reserved.

BIG-IP 2017-10-18 apm profile oauth(1)

apm profile remote-desktop

NAME

remote-desktop - Displays information about a default profile that supports a Citrix remote desktop resource.

MODULE

apm profile

SYNTAX

Displays the properties of the remote-desktop component within the profile module.

DISPLAY

```
list remote-desktop
list remote-desktop [ [ [name] | [glob] | [regex] ] ... ]
show running-config remote-desktop
show running-config remote-desktop [ [glob] | [regex] ] ... ]
options:
  all-properties
  location-specific [true | false]
  non-default-properties
  one-line
```

DESCRIPTION

You can use the remote-desktop component to display the properties of the default remote desktop profile.

A remote desktop profile is for internal use only. You should not create or modify a remote desktop profile.

EXAMPLES

```
list remotedesktop all-properties
```

Displays all of the properties of the default remote desktop profile.

OPTIONS

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

SEE ALSO

itm_virtual

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

apm profile vdi

NAME

vdi - Configures a VDI profile.

MODULE

apm profile

SYNTAX

Configure the vdi component within the profile module using the syntax shown in the following sections.

CREATE/MODIFY

create vdi [name]

modify vdi [name]

options:

msrdp-ntlm-auth-name [[string] | none]

citrix-storefront-replacement [enabled | disabled]

edit vdi [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list vdi

list vdi [[[name] | [glob] | [regex]] ...]

show running-config vdi

show running-config vdi [[glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete vdi [name]

DESCRIPTION

You can use the vdi component to configure a VDI profile. A VDI profile is a group of settings that you can use to enable and configure VDI services such as Citrix, VMware View and MSRDP, so that those work with the BIG-IP system.

EXAMPLES

```
create vdi MyVdiProfile {
  msrdp-ntlm-auth-name "MyNTLMAuth"
  citrix-storefront-replacement disabled
}
```

Creates a VDI profile named MyVdiProfile with NTLM Authentication configuration MyNTLMAuth to be used for MSRDP clients authentication.

```
list vdi all-properties
```

Displays a list of VDI profiles, including parameter values.

```
delete vdi MyVdiProfile
```

Deletes the VDI profile named MyVdiProfile.

OPTIONS

msrdp-ntlm-auth-name

Specifies the NTLM auth configuration object to be used by this VDI profile for MSRDP client authentication.

citrix-storefront-replacement

Specifies whether Citrix StoreFront functionality is enabled on APM. When this option is enabled, APM will respond to StoreFront-specific requests from Citrix Receiver clients. When disabled, APM will fallback to PNAgent protocol.

SEE ALSO

ltm_virtual

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

apm report custom-report-field

NAME

custom-report-field - Configures a custom report field.

MODULE

apm report

SYNTAX

Configure the custom-report-field component within the custom-report module using the syntax shown in the following sections.

CREATE/MODIFY

```
create custom-report-field name [string] field-id [integer] field-position [integer] report-name [string]
```

options

```
alias [string]
app-service [string]
sort-direction [asc|desc|unsorted]
sort-order [integer]
```

```
modify custom-report-field [alias | app-service | field-id | field-position | name | report-name | sort-direction | sort-order]
```

options

```
alias [string]
app-service [string]
field-position [integer]
report-name [string]
sort-direction [asc|desc|unsorted]
sort-order [integer]
```

DISPLAY

```
list custom-report-field
```

```
list custom-report-field [alias | app-service | field-id | field-position | name | report-name | sort-direction | sort-order]
```

DESCRIPTION

Configures a custom report field with options.

It is highly recommended that you use the web GUI to configure this custom object.

EXAMPLES

```
create alias UserName field-id 12 field-position 1 name APDNOTICE_USERNAME.User_Name report-name cust1
sort-direction asc sort-order 100000
```

Creates a custom-report-field named UserName under custom report named cust1, field-position is 1, sort direction is ascending and sort-order is 100000.

Note: Some field-id's and names are predefined and fixed values and not customizable, see below for details:

```
apm report custom-report-field {
field-id 21
name APDNOTICE_CLIENT_INFO.Client_Version
}
apm report custom-report-field {
field-id 19
name APDNOTICE_CLIENT_INFO.Client_UI_Mode
}
apm report custom-report-field {
field-id 16
name APDNOTICE_ASSIGN_WEBTOP_LINKS.Webtop_Links
}
apm report custom-report-field {
field-id 18
name APDNOTICE_ACCESS_POLICY_RESULT.Access_Policy_Result
}
apm report custom-report-field {
field-id 20
name APDNOTICE_CLIENT_INFO.Client_Type
}
apm report custom-report-field {
field-id 15
name APDNOTICE_ASSIGN_WEBTOP.Webtop_Name
}
apm report custom-report-field {
field-id 13
name APDNOTICE_ASSIGN_ACL.ACL_variable
}
```

```

    apm report custom-report-field {
field-id 12
name APDNOTICE_USERNAME.User_Name
    }
    apm report custom-report-field {
field-id 14
name APDNOTICE_ASSIGN_RESOURCE.resource_var
    }
    apm report custom-report-field {
field-id 17
name APDNOTICE_CLIENT_INFO.Client_Platform
    }

```

OPTIONS

alias

Specifies the alias of the custom report field.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

field-id

Specifies a unique id of the custom report field. This item is fixed and not customizable.

field-position

Specifies field's position among all available fields.

name Specifies a unique name for the custom report field. This item is fixed and not customizable.

report-name

Specifies the custom report that the custom-report-field belongs to.

sort-direction

Specifies the sort direction of the custom-report-field. It can be asc, desc, or unsorted. Default is asc.

sort-order

Specifies the sort order of the custom-report-field. Default is 100000.

SEE ALSO

custom-report

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm report custom-report-field(1)

apm resource app-tunnel

NAME

app-tunnel - Configures an application tunnel.

MODULE

apm resource

SYNTAX

Configure the app-tunnel component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

```
create app-tunnel [name]
```

```
modify app-tunnel [name]
```

options:

```
acl-order [integer]
```

```
app-service [[string] | none]
```

```
application-launch-warning [true | false]
```

```
apps [add | delete | modify | replace-all-with] {
[name]
```

```
}
```

```
customization-group [add | delete | modify | replace-all-with] {
[name]
```

```
}
```

```
description [[string] | none]
```

```
location-specific [true | false]
```

type [app-tunnel | last | network-access | remote-desktop | web-application]

edit app-tunnel [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list app-tunnel

list app-tunnel [[[name] | [glob] | [regex]] ...]

show running-config app-tunnel

show running-config app-tunnel [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show app-tunnel

show app-tunnel [name]

DELETE

delete app-tunnel [name]

DESCRIPTION

You can use the app-tunnel component to configure an application tunnel to provide secure access to a network, remote desktop, or specific applications.

EXAMPLES

item create app-tunnel myapptunnel customization-group myapptunnelcg

Creates an application tunnel named myapptunnel that uses the policies in the customization group myapptunnelcg.

item delete app-tunnel myapptunnel

Deletes the application tunnel named myapptunnel.

OPTIONS

acl-order

Specifies the location of this app tunnel in the ACL hierarchy in Access Policy Manager ACL lists. The default is 0 (zero).

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

application-launch-warning

Specifies whether to display a warning before launching an application. The options are:

true The system displays security warnings before launching an application, regardless of whether the site is considered a Trusted site. This is the default value.

false

The system displays security warnings before launching an application, only if the site is not in the Trusted Sites list.

apps Specifies the applications that a user can access using this application tunnel. The default is none.

customization-group

Specifies whether customizations are applied to the application tunnel. You can add, modify, delete, or replace all customization groups. This option is required.

description

Specifies a description for the application tunnel. The default is none.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies a name for the component.

partition

Displays the partition within which the app-tunnel component resides. The default is common.

type Specifies the type of application tunnel. The options are:

app-tunnel

This is the default.

network-access

Provides access to a network.

remote-desktop

Provides access to a remote desktop.

web-application
Provides access to a Web application.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 apm resource app-tunnel(1)

apm resource client-rate-class

NAME

client-rate-class - Configures a client rate class resource.

MODULE

apm resource

SYNTAX

Configure the client-rate-class component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

create client-rate-class [name]

modify client-rate-class [name]

options:

app-service [[string] | none]

burst [integer]

ceiling [integer]

description [[string] | none]

dscp [integer]

location-specific [true | false]

mode [borrow | discard | shape]

rate [integer]

type [best-effort | controlled-load | guaranteed]

edit client-rate-class [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list client-rate-class

list client-rate-class [[[name] | [glob] | [regex]] ...]

show running-config client-rate-class

show running-config client-rate-class [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

show client-rate-class

show client-rate-class [name]

DELETE

delete client-rate-class [name]

DESCRIPTION

You can use the client-rate-class component to configure a client rate class resource, which is used in traffic control.

EXAMPLES

```
create client-rate-class sf1{ dscp 40 rate 60000 ceiling 80000 mode shape }
```

Creates a client rate class resource named sf1 used in traffic control. Sets the dscp to 40 and the rate to 60000, sets the ceiling to 80000, and sets the mode to shape.

```
list client-rate-class all
```

Displays a list of all client rate class on the system.

```
delete client-rate-class sf1
```

Deletes the client rate class named sf1 from the system.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot

modify or delete the object. Only the application service can modify or delete the object.

burst

Specifies in bytes the maximum amount of data that can reach the ceiling rate at one time. The default is 0 (zero).

ceiling

Specifies how far, beyond the value specified for the rate option, that traffic can flow when bursting. This number sets an absolute limit. No traffic can exceed this rate. The rate class might limit traffic throughput to the value of the rate option when there is high contention among siblings of a parent-child class hierarchy. The default value is the value of the rate option. The minimum value is 296 bp.

description

Specifies a description for the client rate class. The default is none.

dscp Specifies six bits of DS field used as a codepoint to select the PHB (Per Hop Behavior) for a packet in each network node. The default is -1.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies a unique name to identify the client rate class.

mode Specifies the mode to use for this client rate class. The options are:

borrow

Allows traffic on the client rate class to borrow resources from other flows that are temporarily idle. Traffic that borrows resources is marked as nonconforming and receives a lower priority. This is the default.

discard

Discards packets that do not conform to the specified traffic control descriptor.

shape

Delays packets submitted for transmission until the packets conform to the specified flow parameters

partition

Displays the partition within which this component resides. The default is common.

rate Specifies the guaranteed throughput rate of the traffic handled by this rate class. You can configure the rate in bits per second (bps), kilobits per second (Kbps), megabits per second (Mbps), or gigabits per second (Gbps).

type Specifies the service type in use for the client rate class. The options are:

best-effort

Windows traffic control creates a flow for this client traffic class, and traffic on the flow is handled with the same priority as other Best Effort traffic. This is the default.

controlled-load

Traffic control transmits a very high percentage of packets to the intended receivers. Packet loss for this type closely approximates the basic packet error rate of the transmission medium. Transmission delay for a very high percentage of the delivered packets does not greatly exceed the minimum transit delay experienced by any successfully delivered packet.

guaranteed

Guarantees that datagrams arrive within a specified delivery time and will not be discarded due to queue overflows, provided that the flow of traffic stays within specified traffic parameters. This type is intended for applications that require guaranteed packet delivery.

SEE ALSO

tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 apm resource client-rate-class(1)

apm resource client-traffic-classifier

NAME

client-traffic-classifier - Configures client traffic classifier entries.

MODULE

apm resource

SYNTAX

Configure the client-traffic-classifier component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

```
create client-traffic-classifier [name]
modify client-traffic-classifier [name]
options:
  app-service [[string] | none]
  entries [add | delete | modify | replace-all-with] {
    [name] {
      app-service [[string] | none]
      client-rate-class [[string] | none]
      dst-ip [[ipv4 address] | none]
      dst-mask [[integer] | none]
      dst-port [[integer] | none]
      protocol [[integer] | none]
      src-ip [[ipv4 address ] | none]
      src-mask [[integer] | none]
      src-port [[integer] | none]
    }
  }
  location-specific [true | false]
}

edit client-traffic-classifier [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list client-traffic-classifier
list client-traffic-classifier [ [ [name] | [glob] | [regex] ] ... ]
show running-config client-traffic-classifier
show running-config client-traffic-classifier [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

```
show client-traffic-classifier
show client-traffic-classifier [name]
```

DELETE

```
delete client-traffic-classifier [name]
```

DESCRIPTION

You can use the client-traffic-classifier component to configure a client traffic classifier, which is used by traffic control agent.

EXAMPLES

```
create client-traffic-classifier tf1 { entries entry1 { protocol "6" dst-ip "192.168.0.0" dst-mask
"255.255.0.0" dst-port "0" client-rate-class "sf1" }
entry2 { protocol "6"
src-ip "10.10.0.0"
src-mask "255.255.255.0"
client-rate-class "sf2" } }
```

Creates a client traffic classifier named tf1, sets the entry to entry1, the protocol to 6, the DST IP to 192.168.0.0, the DST mask to 255.255.0.0, the DST port to 0 (zero), and the client rate class to sf1.

```
list client-traffic-classifier all
Displays a list of all client traffic classifiers on the system.
```

```
modify client-traffic-classifier tf1 entries entry1 protocol 17
Modifies the client traffic classifier named tf1.
```

```
delete client-traffic-classifier tf1
Deletes the client traffic classifier named tf1 from the system.
```

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

dst-ip
Specifies the IP address of the receiver of the packet.

dst-mask
Specifies the subnet mask for the destination address.

dst-port

Specifies the 16-bit number to identify the sending port for either UDP or TCP network application.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies the name of the component.

partition

Displays the partition within which the component resides. The default is Common.

protocol

Specifies which traffic protocol to use in the filtering rule.

src-ip

Specifies the address from which the packet is being sent.

src-mask

Specifies the subnet mask for the source address.

src port

Specifies a 16-bit number to identify the sending port for either UDP or TCP network application.

SEE ALSO

tmssh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 apm resource client-traffic-classifier(1)

apm resource ipv6-leasepool

NAME

ipv6-leasepool - Configures a lease pool.

MODULE

apm resource

SYNTAX

Configure the ipv6-leasepool component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ipv6-leasepool [name]
modify ipv6-leasepool [name]
options
  app-service [[string] | none]
  description [[string] | none]
  location-specific [true | false]
  members [add | delete | modify | replace-all-with] {
    [[first ip address in range] - [last ip address in range]]
  }
```

DISPLAY

```
list ipv6-leasepool
list ipv6-leasepool [ [ [name] | [glob] | [regex] ] ... ]
show running-config ipv6-leasepool
show running-config ipv6-leasepool [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show ipv6-leasepool
```

```
show ipv6-leasepool [name]
```

DELETE

```
delete ipv6-leasepool [name]
```

DESCRIPTION

Configures an IPv6 lease pool to create a collection of IPv6 addresses grouped as a single object. You can use a lease pool to associate that collection of IP addresses with a network access resource.

EXAMPLES

```
create ipv6-leasepool myipv6-leasepool {fd1f::1-fd1f::64}
```

Creates a ipv6-leasepool named myipv6-leasepool that contains the IPv6 addresses in the range fd1f::1 - fd1f::64.

Note: No spaces are allowed between the first IPv6 address, hyphen, and second IPv6 address.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

Specifies a unique description of the lease pool.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies a unique name for the lease pool.

members

Specifies a range of IPv6 addresses separated by a hyphen.

partition

Displays the partition within which the component resides. The default is Common.

SEE ALSO

apm profile, ltm virtual

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-24 apm resource ipv6-leasepool(1)

apm resource leasepool

NAME

leasepool - Configures a lease pool.

MODULE

apm resource

SYNTAX

Configure the leasepool component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

```
create leasepool [name]
```

```
modify leasepool [name]
```

options

```
app-service [[string] | none]
```

```
description [[string] | none]
```

```
location-specific [true | false]
```

```
members [add | delete | modify | replace-all-with] {  
  [[first ip address in range] - [last ip address in range]]  
}
```

DISPLAY

```
list leasepool
```

```
list leasepool [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config leasepool
```

```
show running-config leasepool [ [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties

non-default-properties

one-line
partition

show leasepool
show leasepool [name]

DELETE
delete leasepool [name]

DESCRIPTION

Configures a lease pool to create a collection of IPv4 addresses grouped as a single object. You can use a lease pool to associate that collection of IPv4 addresses with a network access resource.

EXAMPLES

```
create leasepool myleasepool {10.10.10.1-10.10.10.10}
```

Creates a leasepool named myleasepool that contains the IPv4 addresses in the range 10.10.10.1 - 10.10.10.10.

Note: No spaces are allowed between the first IPv4 address, hyphen, and second IPv4 address.

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description
Specifies a unique description of the lease pool.

location-specific
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]
Specifies a unique name for the lease pool.

members
Specifies a range of IP addresses separated by a hyphen.

partition
Displays the partition within which the component resides. The default is Common.

SEE ALSO

apm profile, ltm virtual

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-24 apm resource leasepool(1)

apm resource network-access

NAME
network-access - Configures general settings for a network access connection.

MODULE
apm resource

SYNTAX
Configure the network-access component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY
create network-access [name]
modify network-access [name]
options:
app-service [[string] | none]
address-space-dhcp-requests-excluded [true | false]
address-space-exclude-subnet [[string] | none]
ipv6-address-space-exclude-subnet [[string] | none]
address-space-include-dns-name [[string] | none]
address-space-exclude-dns-name [[string] | none]

address-space-include-subnet [[string] | none]
ipv6-address-space-include-subnet [[string] | none]
address-space-local-subnets-excluded [true | false]
address-space-loc-dns-servers-excluded [true | false]
address-space-protect [true | false]
application-launch [[string] | none]
application-launch-warning [true | false]
auto-launch [true | false]
client-interface-speed [[integer] | none]
client-ip-filter-engine [true | false]
client-power-management [ignore | prevent | terminate]
client-proxy [true | false]
client-proxy-address [ip addr]
client-proxy-enforce-subnets [true | false]
client-proxy-exclusion-list [[string] | none]
client-proxy-ignore-auto-config-error [true | false]
client-proxy-local-bypass [true | false]
client-proxy-port [[integer] | none]
client-proxy-script [[string] | none]
client-proxy-use-http-pac [true | false]
client-proxy-use-local-proxy [true | false]
client-traffic-classifier [[string] | none]
compression [gzip | none]
customization-group [[string] | none]
description [[string] | none]
dns-primary [ip addr]
ipv6-dns-primary [ip addr]
dns-secondary [ip addr]
ipv6-dns-secondary [ip addr]
dns-suffix [[string] | none]
drive-mapping [[string] | none]
dtls [true | false]
dtls-port [[integer] | none]
execute-logoff-scripts [true | false]
idle-timeout-threshold [[integer] | none]
idle-timeout-window [[integer] | none]
leasepool-name [[string] | none]
location-specific [true | false]
ipv6-leasepool-name [[string] | none]
microsoft-network-client [true | false]
microsoft-network-server [true | false]
network-tunnel [enabled | disabled]
optimized-app [add | delete | modify | none | replace-all-with]
provide-client-cert [true | false]
proxy-arp [true | false]
split-tunneling [true | false]
static-host [[string] | none]
supported-ip-version [ipv4 | ipv4-ipv6]
sync-with-active-directory [true | false]
type [app-tunnel | last | network-access | remote-desktop | web-application]
wins-primary [ip addr]
wins-secondary [ip addr]

edit network-access [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list network-access

list network-access [[[name] | [glob] | [regex]] ...]

show running-config network-access

show running-config network-access [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

show network-access

show network-access [name]

DELETE

delete network-access [name]

DESCRIPTION

You can use the network-access component to configure the general settings for a network access connection.

EXAMPLES

create network-access mynetwork-access customization-group mynetaccess

Creates a network access connection configuration object named mynetwork-access that uses the policies in the customization group named mynetaccess.

delete network-access mynetwork-access

Deletes the network access connection configuration object named mynetwork-access.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

address-space-dhcp-requests-excluded

Specifies whether requests from IP addresses using DHCP are excluded from accessing the network. The default is true.

address-space-exclude-subnet

Specifies the IPv4 address spaces whose traffic you want to exclude from access to a subnet on the network. The default is none.

ipv6-address-space-exclude-subnet

Specifies the IPv6 address spaces whose traffic you want to exclude from access to a subnet on the network. The default is none.

address-space-include-dns-name

Specifies a list of domain names describing the target LAN DNS addresses for split tunneling only. You can add multiple address spaces to the list. For each address space, type the domain name, in the form site.siterequest.com or *.siterequest.com. The default is none.

address-space-exclude-dns-name

Specifies the DNS address spaces whose traffic you want to exclude from access to a subnet on the network. You can add multiple address spaces to the list. For each address space, type the domain name, in the form site.siterequest.com or *.siterequest.com. The default is none.

address-space-include-subnet

Specifies a list of IPv4 addresses or address/mask pairs describing the target LAN. When using split tunneling, only the traffic to these addresses and network segments goes through the tunnel configured for Network Access. You can add multiple address spaces to the list. For each address space, type the IPv4 address and network mask. The default is none.

ipv6-address-space-include-subnet

Specifies a list of IPv6 addresses or address/mask pairs describing the target LAN. When using split tunneling, only the traffic to these addresses and network segments goes through the tunnel configured for Network Access. You can add multiple address spaces to the list. For each address space, type the IPv6 address and network mask. The default is none.

address-space-local-subnets-excluded

Specifies whether to exclude local access to any host or subnet in routes that you have specified in the client routing table. The default is false. When you set this option to true, the system does not support integrated IP filtering.

address-space-loc-dns-servers-excluded

Specifies whether to exclude local access to DNS servers configured on client prior to establishing network access connection. The default is false.

address-space-protect

Specifies whether the IP address spaces whose traffic is forced through the tunnel are protected. The default is false.

app-service

The default is none.

application-launch

Specifies the applications to launch when the client accesses the network. The default is none.

application-launch-warning

Specifies whether the user is warned that an application is being launched. The default is true.

auto-launch

Specifies whether NA resource is to be launched automatically from full webtop. The default is false.

client-interface-speed

Specifies the baud rate of the client interface with the network. The default is 100000000.

client-ip-filter-engine

Specifies whether the client IP address is filtered. The default is .

client-power-management

Specifies how to interact with Windows power management features.

prevent

Prevents Windows from entering standby/hibernate during connection.

terminate

Terminate network access connection if Windows is entering standby/hibernate

ignore

Do nothing. Ignore power management events. This is the default value.

client-proxy

Specifies whether this resource handles a client proxy. The default is false.

client-proxy-address

Specifies the IP address of the proxy client. The default is any6.

client-proxy-enforce-subnets

Specifies whether address space subnets must be enforced in proxy auto-configuration. The default is true.

client-proxy-exclusion-list

Specifies the Web addresses that do not need to be accessed through your proxy server. You can use wild cards to match domain and host names or addresses, for example, www.*.com, 128.*, 240.8, 8., mygroup.*, and *.*. The default is none.

client-proxy-ignore-auto-config-error

Allow client to connect even after an error in merging or downloading a proxy auto-configuration file. The default is false.

client-proxy-local-bypass

Specifies whether you want to allow local (intranet) addresses to bypass the proxy server. The default is false.

client-proxy-port

Specifies the port number of the proxy server you want Network Access clients to use to connect to the Internet. The default is 0 (zero).

client-proxy-script

Specifies the URL for a proxy auto-configuration script, if one is used with this connection. The default is none.

client-proxy-use-http-pac

Specifies whether the browser uses http:// to locate the proxy the autoconfig file, instead of file://. Set this to true for applications, like Citrix MetaFrame, that cannot use the client proxy autoconfig script when the browser attempts to use the prefix file:// to locate the script. The default is false.

client-proxy-use-local-proxy

Specifies whether the browser uses the proxy configured on client prior to establishing network access connection. The default is false.

client-traffic-classifier

Specifies a client traffic classifier to use with this network access connection. The default is none.

compression

Specifies whether you want to compress all traffic between the Network Access client and the controller. The default is none.

customization-group

Specifies the customization group that defines the policies that apply to network access. This option is required.

description

Specifies a unique description of the network access configuration object. The default is none.

dns-primary

For split tunneling, specifies the IPv4 address of the primary name server that is conveyed to the remote access point for IPv4 traffic. The default is any6.

ipv6-dns-primary

For split tunneling, specifies the IPv6 address of the primary name server that is conveyed to the remote access point for IPv6 traffic. The default is any6.

dns-secondary

For split tunneling, specifies the IPv4 address of the secondary name server that is conveyed to the remote access point for IPv4 traffic. The default is any6.

ipv6-dns-secondary

For split tunneling, specifies the IPv6 address of the secondary name server that is conveyed to the remote access point for IPv6 traffic. The default is any6.

dns-suffix

Type in a DNS suffix to send to the client. If this field is left blank, the controller sends its own DNS suffix. You can specify multiple default domain suffixes separated with commas. The default is none.

drive-mapping

For split tunneling, specifies the drive to which this resource provides a network access connection. The default is none.

dtls Specifies whether the network access connection uses Datagram Transport Level Security (DTLS). DTLS uses UDP instead of TCP, to provides better throughput for high demand applications like VoIP or streaming video, especially with lossy connections. The default is false.

dtls-port

Specifies the port number that the network access resource uses for secure UDP traffic with DTLS. The default is 4433.

execute-logout-scripts

Specifies whether the system to executes logout scripts (configured on the Active Directory domain) when the connection is terminated. The default is false.

idle-timeout-threshold

Defines the average byte rate that either ingress or egress tunnel traffic must exceed for the tunnel to

update a session. If the average byte rate falls below the specified threshold, the system applies the inactivity timeout, which is defined in the session's Access Profile. The default is 0 (zero).

idle-timeout-window

Defines the value that the system uses to calculate the Exponential Moving Average (EMA) byte rate of ingress and egress tunnel traffic. The default is 0 (zero).

leasepool-name

Specifies the IPv4 lease pools that the user can access with this network access connection. The default is none.

ipv6-leasepool-name

Specifies the IPv6 lease pools that the user can access with this network access connection. The default is none.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

microsoft-network-client

Specifies whether the client PC can access remote resources over a VPN connection. The default is true.

microsoft-network-server

Specifies whether the server can access remote resources over a VPN connection. The default is false.

network-tunnel

Enables or disables the network tunnel. The default is enabled.

optimized-app

Specifies the optimized applications that you want to users to access using this network access connection resource. You can add, delete, modify, or replace the current optimized applications. The default is none.

partition

Displays the partition within which this network access connection component resides. The default is Common.

provide-client-cert

Specifies whether client certificates are required to establish an SSL connection. You can set this option to false if the client certificates are only requested in an SSL connection. In this case, the client is configured to not send client certificates. The default is true.

proxy-arp

Select Enable to enable Proxy ARP for this network access resource. When you implement Proxy ARP for a network access resource, remote VPN tunnel clients can use IP addresses from the LAN IP subnet without additional network infrastructure changes. Ranges of IP addresses from the LAN subnet can be configured in the lease pools and assigned to tunnel clients. When a host on the LAN sends traffic to a tunnel client, an ARP query is sent to request the client address. Access Policy Manager then responds with its own MAC address. Traffic is then sent to network access and forwarded to the client over the network access tunnel. No configuration changes are required on devices other than the Access Policy Manager.

See your Network Access documentation for more information about Proxy ARP configuration. The default is false.

split-tunneling

Specifies whether only traffic targeted to a specified address space is sent over the network access tunnel. With split tunneling, all other traffic bypasses the tunnel. The default is false. When you set this option to true, all traffic passing over the network access connection uses this setting.

static-host

Specifies the static hosts to which this resource provides a network access connection. The default is none.

supported-ip-version

Specifies the supported IP protocol version. The default is ipv4.

sync-with-active-directory

Specifies whether you want the network access connection to emulate the Windows logon process for a client on an Active Directory domain. The default is false.

When this option is set to true, network policies are synchronized when the connection is established, or at logoff. The following items are synchronized:

• Logon scripts are started as specified in the user profile.

• Drives are mapped as specified in the user profile.

• Group policies are synchronized as specified in the user profile. Group Policy logon scripts are started when the connection is established, and Group Policy logoff scripts are run when the network access connection is stopped.

type Specifies the type of network access connection this component provides. The default is network-access.

wins-primary

Specifies the primary IP address to which this resource provides a network access connection. The default

is any6.

wins-secondary

Specifies the secondary IP address to which this resource provides a network access connection. The default is any6.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2017-05-09 apm resource network-access(1)

apm resource portal-access

NAME

portal-access - Configures a portal access resource.

MODULE

apm resource

SYNTAX

Configure the portal-access component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

create portal-access [name]

modify portal-access [name]

options:

acl-order [integer]

application-uri [string] | none]

app-service [[string] | none]

css-patching [true | false]

customization-group [string] | none]

description [string] | none]

flash-patching [true | false]

host-replace-string [string] | none]

host-search-strings [string] | none]

html-patching [true | false]

items [add | delete | modify | replace-all-with] {
[string]

}

javascript-patching [true | false]

location-specific [true | false]

patching-type [full-patch | min-patch]

path-match-case [true | false]

proxy-host [string] | none]

proxy-port [string] | none]

publish-on-webtop [true | false]

scheme-patching [true | false]

edit portal-access [all-properties | non-default-properties]

options:

all-properties

non-default-properties

DISPLAY

list portal-access

list portal-access [[[name] | [glob] | [regex]] ...]

show running-config portal-access

show running-config portal-access [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show portal-access

show portal-access [name]

DELETE

delete portal-access [name]

DESCRIPTION

You can use the portal-access component to specify a portal access resource.

EXAMPLES

```
item create portal-access myportalaccess acl-order 14 patching-type full-patch items add { item1 { host www.mywebsite.com paths /* } }
```

Creates a portal access resource named myportalaccess.

```
item delete portal-access myportalaccess
```

Deletes the portal access resource named myportalaccess.

OPTIONS

`acl-order`
Specifies the order of this portal access in Access Policy Manager ACL lists. This option is required.

`application-uri`
`app-service`
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`css-patching`
Specifies whether the response content type CSS is patched. The default is true.

`customization-group`
The customization group is created automatically if not specified.

`description`
Specifies a description of the resource. The default is none.

`flash-patching`
Specifies whether the system patches Flash content. The default is true.

`host-replace-string`
Specifies the replacement host string, when you specify minimal for the patching-type option.

`host-search-strings`
Specifies the host string to replace, when you specify minimal for the patching-type option.

`html-patching`
Specifies whether the system patches HTML content. The default is true.

`items`
Configures the host name or IP address, the network mask (if the resource is a network), the port, and any paths specified for a portal access resource. The default is none.

`javascript-patching`
Specifies whether the system patches JavaScript content. The default is true.

`location-specific`
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

`[name]`
Specifies a unique name for the component.

`patching-type`
Specifies whether this resource provides minimal or full path patching.

`path-match-case`
Specifies whether the application URI is case-sensitive. The default is true.

`proxy-host`
Specifies the proxy host that the portal access uses. The default is none. If you configure this option, you must also configure the option proxy-port.

`proxy-port`
Specifies the port that the portal access proxy uses. The default is none. Configure this option, only when you configure the option proxy-host.

`publish-on-webtop`
Specifies whether to publish this resource on the webtop. The default is false. If you set this option to true, you must also specify the Application URI using the application-uri option.

`scheme-patching`
Specifies whether this resource replaces all HTTP scheme addresses with HTTPS scheme addresses. This option is effective only when minimal patching is selected for patching-type. The default is false.

SEE ALSO

tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-12-03 apm resource portal-access(1)

apm resource remote-desktop citrix-client-bundle

NAME

citrix-client-bundle - Configures a Citrix Client Bundle remote desktop resource configuration object.

MODULE

apm resource remote-desktop

SYNTAX

Configure the citrix-client-bundle component within the resource remote desktop module using the syntax shown in the following sections.

CREATE/MODIFY

create citrix-client-bundle [name]

modify citrix-client-bundle [name]

options:

app-service [[string] | none]

download-url [[url] | none]

packages [[string] | none]

windows-download-url [[url] | none]

windows-min-version [[string] | none]

windows-package [[string] | none]

sb-windows-package [[string] | none]

edit citrix-client-bundle [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list citrix-client-bundle

list citrix-client-bundle [[[name] | [glob] | [regex]] ...]

show running-config citrix-client-bundle

show running-config citrix-client-bundle [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show citrix-client-bundle

show citrix-client-bundle [name]

DELETE

delete citrix-client-bundle [name]

DESCRIPTION

You can use the citrix-client-bundle component to configure a Citrix Client Bundle remote desktop resource.

EXAMPLES

```
create citrix-client_bundle myccb { windows-min-version xp }
```

Creates a Citrix Client Bundle remote desktop resource named myccb that can be downloaded from receiver.citrix.com (the default value), where the client must have at least Windows XP installed.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

download-url

Specifies the default location receiver.citrix.com from which to download the Citrix installation package.

packages

This option is deprecated and is maintained here for backward compatibility reasons.

[name]

Specifies an object name. This option is required; however, the parameter name is implicit and must not be typed in the syntax.

windows-download-url

Specifies the location from which to download the Windows version. You can provide a value for either the windows-download-url or windows-package option, but not both. The default is none.

windows-min-version

Specifies the oldest version of the Citrix client that can be used with this remote desktop resource. The default is none.

windows-package

This option is deprecated and is replaced by sb-windows-package and is maintained here for backward compatibility reasons.

sb-windows-package

Specifies the location of the Citrix Receiver client or Citrix HTML5 package to be uploaded. You can provide a value for either the sb-windows-package or windows-download-url option, but not for both. The default is none. The package can be uploaded via BIG-IP TMUI or can be assigned with TMSH. Follow the below mentioned steps to assign package to sb-windows-package.

Configure the component of the sandbox, citrix-client-package, within the resource module under apm using the syntax shown in the following sections.

MODIFY

```
modify sandbox citrix-client-package
options
  base-uri [string]
  description [[string] | none]
  files [add | delete | modify | replace-all-with] {
    [item name] {
  content-type [string]
  filename [string]
  file-type citrix-bundle
  folder [string]
  local-path [string]
  name [string] } }
```

DISPLAY

```
list sandbox
list sandbox [ [name] | [glob] | [regex] ] ... ]
  all-properties
  non-default-properties
  one-line
  partition
```

A sandbox is a container for files stored on the BIG-IP, to which you want to provide client access. The local-path indicates the location of the file on the BIG-IP to be added into the sandbox e.g. /shared/tmp/FILENAME. Once the item is created on modifying citrix-client-package sandbox, it can be referenced as citrix-client-package:[item name] while assigning it to sb-windows-package.

SEE ALSO

citrix, citrix-client-package-file, rdp

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2017-11-21 apm resource remote-desktop citrix-client-bundle(1)

apm resource remote-desktop citrix-client-package-file

NAME

citrix-client-package-file - Configures a Citrix client package file configuration object.

MODULE

apm resource remote-desktop

SYNTAX

Configure the citrix-client-package-file component within the resource remote desktop module using the syntax shown in the following sections.

CREATE/MODIFY

```
create citrix-client-package-file [name]
modify citrix-client-package-file [name]
options:
  app-service [[string] | none]
  location-specific [true | false]
  original-file-name [[string] | none]
  source-path [[string] | none]
```

edit citrix-client-package-file [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties

DISPLAY

list citrix-client-package-file

list citrix-client-package-file [[[name] | [glob] | [regex]] ...]

show running-config citrix-client-package-file

show running-config citrix-client-package-file [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line
- partition

DELETE

delete citrix-client-package-file [name]

DESCRIPTION

You can use the citrix-client-package-file component to configure access to a Citrix client package file.

EXAMPLES

```
create citrix-client-package myccpackage { source-path www.siterequest.citrix_download.com }
```

Creates a Citrix client package remote desktop resource named myccpackage that is available from www.siterequest.citrix_download.com.

OPTIONS

[name]

Specifies an object name. This option is required; however, the parameter name is implicit and must not be typed in the syntax.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

original-file-name

Specifies the original file name of the Citrix Installation package file name to download. The default is none.

source-path

Specifies the location from which to download the Citrix client package file. This option is required.

SEE ALSO

citrix, citrix-client-bundle, rdp

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-1apm4resource remote-desktop citrix-client-package-file(1)

apm resource remote-desktop citrix

NAME

citrix - Configures a Citrix remote desktop resource configuration object.

MODULE

apm resource remote-desktop

SYNTAX

Configure the citrix component within the resource remote desktop module using the syntax shown in the following sections.

CREATE/MODIFY

create citrix [name]

modify citrix [name]

options:

```
app-service [[string] | none]
auto-logon [enabled | disabled]
customization-group [add | delete | modify | replace-all-with] {
  [name] {
options:
caption [[string] | none]
detailed-description [[string] | none]
  }
}
description [[string] | none]
domain-source [session.logon.last.domain | none]
enable-serverside-ssl [enabled | disabled]
pool [pool name]
host [fqdn]
ip [ip address]
location-specific [true | false]
password-source [session.logon.last.password | none]
port [[string] | none]
username-source [session.logon.last.username | none]
```

```
edit citrix [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
DISPLAY
list citrix
list citrix [ [ [name] | [glob] | [regex] ] ... ]
show running-config citrix
show running-config citrix [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
show citrix
show citrix [name]
```

```
DELETE
delete citrix [name]
```

DESCRIPTION
You can use the B component to configure a Citrix remote desktop resource.

EXAMPLES

```
create citrix mycitrix { ip 172.29.67.130 }
```

Creates a Citrix remote desktop resource named mycitrix with Citrix XML Broker server specified as IP address 172.29.67.130.

```
create citrix mycitrix { host mycitrix.mycompany.com auto-logon enabled }
```

Creates a Citrix resource with Citrix XML Broker server specified as hostname mycitrix.mycompany.com and auto-logon enabled with APM credentials (that user types on Logon Page).

```
create citrix mycitrix { pool /Common/mycitrix-pool enable-serverside-ssl enabled }
```

Creates a Citrix resource with Citrix XML Broker server(s) specified in pool named /Common/mycitrix-pool and SSL communication enabled to the server(s) (SSL should also be enabled on the servers and APM virtual should have serversssl profile).

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auto-logon
Enables or disables automatic log on to the Citrix server. If you enable this option, you must also provide values for the username-source, password-source, and domain-source options. The default is disabled.

customization-group
Specifies whether customization groups are applied to the Citrix remote desktop. You can add, modify, or delete customization groups. You can also replace all current customization groups with new customization groups. The default is none.

description
Specifies a description for your Citrix remote desktop. The default is none.

domain-source
Specifies the Session variable used as a source for the auto-logon user password. The default is session.logon.last.domain.

enable-serverside-ssl
Enables or disables SSL capabilities between the BIG-IP system and the Citrix server. When enabled, the port number automatically changes to 443. The default is disabled.

pool Specifies the pool name that contains your Citrix XML Broker server(s). You must use one of these options to specify the server address: pool, host, or ip.

`host` Specifies the hostname of your Citrix XML Broker server. You must use one of these options to specify the server address: `pool`, `host`, or `ip`.

`ip` Specifies the IP address of your Citrix XML Broker server. You must use one of these options to specify the server address: `pool`, `host`, or `ip`.

`location-specific`
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

`[name]`
Specifies an object name. This option is required; however, the parameter name is implicit and must not be typed in the syntax.

`password-source`
Specifies the session variable that is used as a source for the auto-logon password. The default is `session.logon.last.password`.

`port` Specifies the port for your Citrix server. The default is 80.

`username-source`
Specifies the session variable that is used as a source for the auto-logon user name. The default is `session.logon.last.username`.

SEE ALSO

`citrix-client-bundle`, `citrix-client-package-file`, `rdp`, `vmware-view`, `quest`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-06-11 apm resource remote-desktop citrix(1)

apm resource remote-desktop quest

NAME

`quest` - Configures a Quest vWorkspace remote desktop resource configuration object.

MODULE

`apm resource remote-desktop`

SYNTAX

Configure the `quest` component within the resource remote desktop module using the syntax shown in the following sections.

CREATE/MODIFY

```
create quest [name]
```

```
modify quest [name]
```

options:

```
app-service [[string] | none]
```

```
auto-logon [enabled | disabled]
```

```
customization-group [add | delete | modify | replace-all-with] {  
  [name] {
```

```
options:
```

```
caption [[string] | none]
```

```
detailed-description [[string] | none]
```

```
}
```

```
description [[string] | none]
```

```
domain-source [session.logon.last.domain | none]
```

```
enable-serverside-ssl [enabled | disabled]
```

```
pool [pool name]
```

```
host [fqdn]
```

```
ip [ip address]
```

```
location-specific [true | false]
```

```
password-source [session.logon.last.password | none]
```

```
port [[string] | none]
```

```
username-source [session.logon.last.username | none]
```

```
edit quest [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY
list quest
list quest [[[name] | [glob] | [regex]] ...]
show running-config quest
show running-config quest [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

show quest
show quest [name]

DELETE
delete quest [name]

DESCRIPTION
You can use the B component to configure a Quest vWorkspace remote desktop resource.

EXAMPLES
create quest myquest { ip 172.29.67.130 }
Creates a Quest vWorkspace remote desktop resource named myquest with the Quest vWorkspace connection broker server specified as IP address 172.29.67.130.

create quest myquest { host myquest.mycompany.com auto-logon enabled }
Creates a Quest vWorkspace resource with the Quest vWorkspace connection broker server specified as hostname myquest.mycompany.com and with auto-logon enabled using the credentials that the user types into the access policy Logon Page.

create quest myquest { pool /Common/myquest-pool enable-serverside-ssl enabled }
Creates a Quest vWorkspace resource with the Quest vWorkspace connection broker servers specified in a pool named /Common/myquest-pool and with SSL communication enabled to the servers. Note: SSL should also be enabled on the servers themselves and the APM virtual server should specify a server SSL profile.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auto-logon
Enables or disables automatic log on to the Quest vWorkspace connection broker server. If you enable this option, you must also provide values for the username-source, password-source, and domain-source options. The default is disabled.

customization-group
Specifies whether customization groups are applied to the Quest vWorkspace resource. You can add, modify, or delete customization groups. You can also replace all current customization groups with new customization groups. The default is none.

description
Specifies a description for your Quest vWorkspace remote desktop. The default is none.

domain-source
Specifies the session variable to use as a source for the auto-logon user password. The default is session.logon.last.domain.

enable-serverside-ssl
Enables or disables SSL capabilities between the BIG-IP system and the Quest vWorkspace connection broker server. When enabled, the port number automatically changes to 443. The default is disabled.

pool Specifies the pool name that contains your Quest vWorkspace connection broker servers. (You must specify the server address using one of these options: pool, host, or ip.)

host Specifies the hostname of your Quest vWorkspace connection broker server. (You must specify the server address using one of these options: pool, host, or ip.)

ip Specifies the IP address of your Quest vWorkspace connection broker server. (You must specify the server address using one of these options: pool, host, or ip.)

location-specific
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]
Specifies an object name. This option is required. Note: The parameter name is implicit. Do not type name in the syntax.

password-source
Specifies the session variable to use as a source for the auto-logon password. The default is session.logon.last.password.

port Specifies the port for your Quest vWorkspace connection broker server. The default is 8080.

username-source
Specifies the session variable to use as a source for the auto-logon user name. The default is

session.logon.last.username.

SEE ALSO

citrix, rdp, vmware-view

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-06-17 apm resource remote-desktop quest(1)

apm resource remote-desktop rdp

NAME

rdp - Configures a Microsoft Remote Desktop Protocol (MSRDP) configuration object.

MODULE

apm resource remote-desktop

SYNTAX

Configure the rdp component within the resource remote desktop module using the syntax shown in the following sections.

CREATE/MODIFY

create rdp [name]

modify rdp [name]

options:

app-service [[string] | none]

auto-logon [enabled | disabled]

customization-group [add | delete | modify | replace-all-with] {
[name] }

options:

app-service [[string] | none]

caption [[string] | none]

detailed-description [[string] | none]

}

}

description [[string] | none]

domain-source [session.logon.last.domain | none]

host [fqdn]

ip [ip address]

location-specific [true | false]

password-source [session.logon.last.password | none]

port [[integer] | none]

username-source [session.logon.last.username | none]

edit rdp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list rdp

list rdp [[[name] | [glob] | [regex]] ...]

show running-config rdp

show running-config rdp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

show rdp

show rdp [name]

DELETE

delete rdp [name]

DESCRIPTION

You can use the rdp component to configure an MSRDP resource.

EXAMPLES

```
create rdp myrdp { host 172.29.67.130 }
```

Creates a MSRDP remote desktop resource named myrdp with an MSRDP server with an IP address of 172.29.67.130.

```
create rdp myrdp { host 172.29.67.130 rdp-cache-bitmaps true }
```

Creates a MSRDP remote desktop resource named myrdp with an MSRDP server with an IP address of 172.29.67.130 where bitmaps are cached on the client PC.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auto-logon

Specifies if automatic log on to the Microsoft RDP server is used. If you enable this option, you must also provide values for the username-source, password-source, and domain-source options. The default is disabled.

customization-group

Specifies whether customization-groups are applied to the remote desktop. You can add, modify, delete, or replace all customization-groups. The default is none.

description

Specifies a description of an MSRDP resource. The default is none.

domain-source

Specifies the session variable used as a source for the auto-logon user password. The default is session.logon.last.domain.

host Specifies the hostname of your Microsoft RDP server. Either the host or ip option is required; however, you cannot specify both options.

ip Specifies the IP address of your Microsoft RDP server. Either the host or ip option is required; however, you cannot specify both options.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies an object name. This option is required; however, the parameter name is implicit and must not be typed in the syntax.

password-source

Specifies the session variable used as a source for the auto-logon password. The default is session.logon.last.password.

port Specify port 3389 for your Microsoft RDP server. The default is 0 (zero).

username-source

Specifies the session variable used as a source for the auto-logon user name. The default is session.logon.last.username.

SEE ALSO

citrix, citrix-client-bundle, citrix-client-package-file, vmware-view, quest

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2019-08-19 apm resource remote-desktop rdp(1)

apm resource remote-desktop vmware-view

NAME

vmware-view - Configures a VMware View remote desktop resource configuration object.

MODULE

apm resource remote-desktop

SYNTAX

Configure the vmware-view component within the resource remote desktop module using the syntax shown in the following sections.

CREATE/MODIFY

create vmware-view [name]

modify vmware-view [name]

options:

app-service [[string] | none]

auto-logon [enabled | disabled]

```
customization-group [add | delete | modify | replace-all-with] {
  [name] {
options:
caption [[string] | none]
detailed-description [[string] | none]
  }
}
description [[string] | none]
domain-source [session.logon.last.domain | none]
enable-serverside-ssl [enabled | disabled]
pool [pool name]
host [fqdn]
ip [ip address]
location-specific [true | false]
password-source [session.logon.last.password | none]
port [[string] | none]
username-source [session.logon.last.username | none]
```

```
edit vmware-view [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
DISPLAY
list vmware-view
list vmware-view [ [ [name] | [glob] | [regex] ] ... ]
show running-config vmware-view
show running-config vmware-view [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
show vmware-view
show vmware-view [name]
```

```
DELETE
delete vmware-view [name]
```

DESCRIPTION

You can use the B component to configure a VMware View remote desktop resource.

EXAMPLES

```
create vmware-view myview { ip 172.29.67.130 }
Creates a VMware View remote desktop resource named myview with the VMware View Connection server specified as IP address 172.29.67.130.
```

```
create vmware-view myview { host myview.mycompany.com auto-logon enabled }
Creates a VMware View resource with the VMware View Connection server specified as hostname myview.mycompany.com and auto-logon enabled with APM credentials (that user types on Logon Page).
```

```
create vmware-view mview { pool /Common/myview-pool enable-serverside-ssl enabled }
Creates a VMware View resource with the VMware View Connection server(s) specified in pool named /Common/myview-pool and SSL communication enabled to the server(s) (SSL should also be enabled on the servers and APM virtual should have serverssl profile).
```

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auto-logon
Enables or disables automatic log on to the VMware View Connection Server server. If you enable this option, you must also provide values for the username-source, password-source, and domain-source options. The default is disabled.

customization-group
Specifies whether customization groups are applied to the VMware View resource. You can add, modify, or delete customization groups. You can also replace all current customization groups with new customization groups. The default is none.

description
Specifies a description for your VMware View remote desktop. The default is none.

domain-source
Specifies the Session variable used as a source for the auto-logon user password. The default is session.logon.last.domain.

enable-serverside-ssl
Enables or disables SSL capabilities between the BIG-IP system and the VMware View Connection server. When enabled, the port number automatically changes to 443. The default is disabled.

pool Specifies the pool name that contains your VMware View Connection server(s). You must use one of these options to specify the server address: pool, host, or ip.

host Specifies the hostname of your VMware View Connection server. You must use one of these options to specify the server address: pool, host, or ip.

ip Specifies the IP address of your VMware View Connection server. You must use one of these options to specify the server address: pool, host, or ip.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies an object name. This option is required; however, the parameter name is implicit and must not be typed in the syntax.

password-source

Specifies the session variable that is used as a source for the auto-logon password. The default is session.logon.last.password.

port Specifies the port for your VMware View Connection server. The default is 80.

username-source

Specifies the session variable that is used as a source for the auto-logon user name. The default is session.logon.last.username.

SEE ALSO

citrix, rdp, quest

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-06-11 apm resource remote-desktop vmware-view(1)

apm resource sandbox

NAME

sandbox - Configures a sandbox.

MODULE

apm resource

SYNTAX

Configure the sandbox component within the resource module using the syntax shown in the following sections.

CREATE

The CREATE command is currently not available. However, a number of sandboxes have already been created. Use these to upload files.

MODIFY

```
modify sandbox [name]
options
  base-uri [string]
  description [[string] | none]
  files [add | delete | modify | replace-all-with] {
    [item name] {
  content-type [string]
  filename [string]
  file-type [citrix-bundle | customization | unknown]
  folder [string]
  local-path [string]
  name [string]
    }
  }
}
```

DISPLAY

```
list sandbox
list sandbox [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DESCRIPTION

Configures a sandbox and its files. A sandbox is a container for files stored on the BIG-IP, to which you want to provide client access.

EXAMPLES

```
modify sandbox hosted-content files add { BIGIPEdgeClient.exe { folder /client local-path /tmp/BIGIPEdgeClient.exe } }
```

Adds a file called BIGIPEdgeClient.exe to sandbox named hosted-content. The virtual path to this file consists of the sandbox's base-uri, the file's folder, and the name of the file. Putting these components together, the virtual path for the uploaded file is /public/share/client/BIGIPEdgeClient.exe, where /public/share is the base-uri, /client is the folder, and BIGIPEdgeClient.exe is the filename. The local-path indicates the location of the file on the disk drive to be added into the sandbox.

Note: The file you add must already be on the BIG-IP system.

OPTIONS

base-uri

Specifies the first component of the virtual path to the sandbox file. The base-uri for sandbox "hosted-content" is /public/share. The virtual path to a sandbox box file is made up of three components: base-uri/folder/filename

All files in a sandbox share the same base-uri, but the folder can be different for each file.

description

Specifies a unique description about the sandbox.

files

Specifies the list of files in the sandbox.

item name

Specifies the name of an item in the list of files. You can use the original filename as the item name. Each item name in a sandbox must be unique.

content-type

Specifies the content-type field in a HTTP header such as "image/gif" or "text/plain". If none is provided, tmsh will try its best to provide this value.

filename

Specifies the last component of the virtual path to the sandbox file. We recommend that you use the filename of the original file for this name.

file-type

Specifies the F5 file type. Currently there are only three types: unknown, citrix-bundle, and customization. No value is required if a file is uploaded to sandbox for "citrix-client-bundle", since this sandbox is the repositories for F5 specific type of file. However, for files uploaded to sandbox "hosted-content" if no value is provided, the file type defaults to "unknown".

folder

Specifies the second component of the virtual path to the sandbox file.

local-path

Specifies the location of the file to be inserted into the sandbox. This file must be on the BIG-IP already.

name Specifies a value for the underlying file object. Use this only if you are trying to add more than one sandbox file in a modify command. Otherwise, don't specify a value for this attribute. The value must be specified as follows: full path of sandbox name:item name. For example, if the sandbox name is '/Common/hosted-content' and the item name is 'index.html', the value should be '/Common/hosted-content:index.html'.

SEE ALSO

webtop, webtop-link

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 apm resource sandbox(1)

apm resource webtop-link

NAME

webtop-link - Configures a webtop link resource.

MODULE

apm resource

SYNTAX

Configure the webtop-link component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

```
create webtop-link [name]
modify webtop-link [name]
options:
  application-uri [string]
  app-service [[string] | none]
  customization-group [string]
  description [[string] | none]
  location-specific [true | false]
```

```
edit webtop-link [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
all
```

DISPLAY

```
list webtop-link
list webtop-link [ [name] | [glob] | [regex] ] ... ]
show running-config webtop-link
show running-config webtop-link [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
all-properties
non-default-properties
one-line
```

```
show webtop-link
show webtop-link [name]
```

DELETE

```
delete webtop-link [name]
```

DESCRIPTION

Configures the settings necessary to define a link to a webtop that is displayed to the end-user as part of the access policy execution.

EXAMPLES

```
create webtop-link mywebtoplinkcg1 application-uri "http://www.externalsite.com/"
Creates a webtop named mywebtoplinkcg1 with the application-uri of http://www.externalsite.com/.
```

OPTIONS

application-uri

Specifies the application URI of the external portal to which this resource provides access for this webtop link. This is a required setting.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

customization-group

Specifies the customization settings for the webtop.

Note: You must create a customization group of type webtop before you can create a webtop resource. If you do not specify a customization group, a group will be created automatically.

description

Specifies a description of the resource. The default is none.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

SEE ALSO

tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2012-10-24 apm resource webtop-link(1)

NAME

webtop - Configures a webtop resource.

MODULE

apm resource

SYNTAX

Configure the webtop component within the resource module using the syntax shown in the following sections.

CREATE/MODIFY

create webtop [name]

modify webtop [name]

options:

app-service [[string] | none]

customization-group [string]

description [[string] | none]

location-specific [true | false]

minimize-to-tray [false | true]

portal-access-start-uri [[string] | none]

webtop-type [full | last | network-access | portal-access]

warn-when-closed [false | true]

edit webtop [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list webtop

list webtop [[[name] | [glob] | [regex]] ...]

show running-config webtop

show running-config webtop [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show webtop

show webtop [name]

DELETE

delete webtop [name]

DESCRIPTION

Configures the settings necessary to define the webtop assigned to the end-user as part of the access policy execution.

EXAMPLES

```
create webtop mynawebtop { customization-group mywebtopcg1 minimize-to-tray false }
```

Creates a webtop named mynawebtop with the customization group mywebtopcg1 and the network access minimize-to-tray option set to false.

```
create webtop mywawebtop { customization-group mywebtopcg1 portal-access-start-uri  
"http://www.siterequest.com" }
```

Creates a webtop named mywawebtop with the customization group mywebtopcg1 and the starting URI for the portal access of http://www.siterequest.com.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

customization-group

Specifies the customization settings for the webtop.

Note: You must create a customization group of type webtop before you can create a webtop resource. This option is required.

description

Specifies a description of the resource. The default is none.

portal-access-start-uri

Specifies the URI that the webtop starts. You can only configure this option if you have configured the webtop-type option for portal-access.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

minimize-to-tray

Specifies whether the network access window (launched from the full webtop) is minimized to the system tray automatically after the network access connection starts. The default is true.

You can configure this option only if you configured the webtop-type option as network-access or full.

With a network access webtop, the webtop automatically minimizes to the tray. With a full webtop, the webtop minimizes to the system tray only after the network access connection is started.

warn-when-closed
Specifies whether the network access window (launched from the full webtop) should display a warning message when the webtop closes.

You can configure this option only if you configured the webtop-type option as full.

webtop-type
Specifies the type of webtop this resource creates. The options are:

full A webtop to which you assign a single network access resource, multiple portal access resources, and multiple application access app tunnel resources, or any combination of the three types. This is the default.

last network-access
A webtop to which you assign only a single network access resource.

portal-access
A webtop to which you assign only portal access resources.

SEE ALSO
tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2014-01-21 apm resource webtop(1)

apm saml artifact-resolution-service

NAME
artifact-resolution-service - Specify service used to resolve SAML artifacts

MODULE
apm saml

SYNTAX
Configure the artifact-resolution-service component within the saml module using the syntax shown in the following sections.

CREATE/MODIFY
create artifact-resolution-service [name]
modify artifact-resolution-service [name]
options:
app-service [[string] | none]
artifact-resolution-service-host [[string] | none]
artifact-resolution-service-port [integer]
artifact-send-method [http-post | http-redirect]
artifact-validity [integer]
basic-auth-password [[string] | none]
basic-auth-username [[string] | none]
description [[string] | none]
location-specific [true | false]
virtual-server-name [name]
want-artifact-resolution-rq-signed [true | false]

edit artifact-resolution-service [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list artifact-resolution-service
list artifact-resolution-service [[[name] | [glob] | [regex]] ...]
show running-config artifact-resolution-service
show running-config artifact-resolution-service [[[name] | [glob] | [regex]] ...]
options:
all-properties
app-service
non-default-properties
one-line
partition

DELETE
delete artifact-resolution-service [name]

DESCRIPTION
You can use the artifact-resolution-service to create and manage artifact resolution services.

EXAMPLES
create artifact-resolution-service my_ars {virtual-server-name my_virt}
Creates a SAML artifact resolution service named my_ars. In this example, the virtual server my_virt will be used to receive artifact resolve requests and send artifact responses.

create artifact-resolution-service my_ars1 {virtual-server-name my_virt1 artifact-resolution-service-host bigip.mycompany.com basic-auth-username user basic-auth-password password artifact-send-method http-redirect }
Creates a SAML artifact resolution service named my_ars1. The service requires that artifact resolve requests be sent using the http-redirect method with an authorization header that contains the specified credentials.

list artifact-resolution-service
Displays a list of artifact resolution services.

delete artifact-resolution-service my_ars
Deletes the my_ars artifact resolution service.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

artifact-resolution-service-host
Specifies the hostname of the artifact resolution service.

artifact-resolution-service-port
Specifies the port that artifact resolution service will be listening on.

artifact-send-method
Specifies method resolver will use when sending artifact resolve requests. Default value is http-redirect.

artifact-validity
Specifies in seconds how long an artifact remains valid. Default value is 60 seconds.

basic-auth-password
Specifies the basic authentication password to send with an artifact resolve request to this BIG-IP.

basic-auth-username
Specifies the basic authentication username to send with an artifact resolve request to this BIG-IP.

description
Specifies a unique description for the artifact resolution service. Default is none.

location-specific
Objects of this class might have location-specific attributes. If the object is location-specific, set to true.

virtual-server-name
Specifies the virtual server to be used by the artifact resolution service.

want-artifact-resolution-rq-signed
Specifies whether this BIG-IP requires artifact resolution requests to be signed. Default value is true.

SEE ALSO
COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012, 2014. All rights reserved.

BIG-IP 2014-11-10 apm saml artifact-resolution-service(1)

apm saml attribute-consuming-service

NAME
attribute-consuming-service - Configure a list of SAML attribute consuming services.

MODULE
apm saml

SYNTAX

Configure the attribute-consuming-service component within the saml module using the syntax shown in the following sections.

MODIFY

```
create attribute-consuming-service [name]
modify attribute-consuming-service [name]
options:
  app-service [[string] | none]
  attributes [add | delete | modify | none | replace-all-with] {
    name [string] {
      app-service [[string] | none]
      attribute-name [string]
      friendly-name [string]
      is-required [bool]
      name-format [string]
    }
  }
  service-description [[string] | none]
  service-name [[string] | none]

edit attribute-consuming-service [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list attribute-consuming-service
list attribute-consuming-service [ [ [name] | [glob] | [regex] ] ... ]
show running-config attribute-consuming-service
show running-config attribute-consuming-service [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete attribute-consuming-service [name]
```

DESCRIPTION

You can use attribute-consuming-service to create and manage SAML attribute consuming services. Each attribute consuming service contains a list of SAML attributes. Each attribute consuming service must contain at least one attribute.

EXAMPLES

```
create attribute-consuming-service my_attr_consuming_service attributes add { my_attr { attribute-name
"urn:oid:1.3.6.1.4.1.5923.1.1.1.7" is-required false friendly-name eduPersonEntitlement name-format
"urn:oasis:names:tc:SAML:2.0:attrname-format:uri" } } service-name "Academic Journals R US" service-description
"Academic Journals R US"
```

Creates a new attribute consuming service named my_attr_consuming_service with one attribute named my_attr.

```
list attribute-consuming-service
```

Displays a list of attribute consuming services.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

attributes

Specifies a list of attributes. attribute-name must be unique within an attribute-consuming-service object. friendly-name is a more human-readable form of the attribute name for cases where attribute-name is complex. is-required indicates if the service requires the attribute in order to function at all. name-format is a URI reference representing the classification of the attribute name for purposes of interpreting the name.

service-description

Specifies a description for the attribute consuming service.

service-name

Specifies a unique name for the attribute consuming service.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

apm saml auth-context-class-list

NAME

auth-context-class-list - Configure a list of SAML authentication context classes.

MODULE

apm saml

SYNTAX

Configure the auth-context-class-list component within the saml module using the syntax shown in the following sections.

MODIFY

create auth-context-class-list [name]

modify auth-context-class-list [name]

options:

app-service [[string] | none]

classes [add | delete | modify | none | replace-all-with] {

 name [string] {

 order [integer]

 value [string]

 }

}

description [[string] | none]

edit auth-context-class-list [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list auth-context-class-list

list auth-context-class-list [[[name] | [glob] | [regex]] ...]

show running-config auth-context-class-list

show running-config auth-context-class-list [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete auth-context-class-list [name]

DESCRIPTION

You can use the auth-context-class-list to create and manage lists of SAML authentication context classes.

Each class in the list must contain a unique order and a unique value. Order indicates the relative level of security ranging from 1 (least secure) to 255 (most secure).

EXAMPLES

```
create sp_authn_ctx_classes_list classes add { password { order 1 value
urn:oasis:names:tc:SAML:2.0:ac:classes:Password } kerberos { order 2 value
urn:oasis:names:tc:SAML:2.0:ac:classes:Kerberos } }
```

Creates a new list named 'sp_authn_ctx_classes_list' with two authentication context classes: password and kerberos. Higher order number implies higher security associated with class. In this example, the fact that the kerberos class order is 2 implies that it has higher security than the password class with order 1.

```
modify authentication_contexts_list classes add { SmartcardPKI { order 8 value
urn:oasis:names:tc:SAML:2.0:ac:classes:SmartcardPKI } }
```

Modifies default list of authentication context classes to include a class 'SmartcardPKI' with priority order '8' and value 'urn:oasis:names:tc:SAML:2.0:ac:classes:SmartcardPKI'.

```
modify authentication_contexts_list classes delete { smartcard }
```

Removes authentication context class 'smartcard' from the default list of authentication context classes 'authentication_contexts_list'.

```
list auth-context-class-list
```

Displays default list of authentication context classes.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

classes

Specifies a list of authentication context classes. Properties 'order' and 'value' must be unique within

the auth-context-class-list object. Property 'order' specifies the security of the class in the context of the BIG-IP system; order ranges from the least secure '1' to the most secure '255'. Property 'value' specifies a URL of authentication context class, for example, 'urn:oasis:names:tc:SAML:2.0:ac:classes:Kerberos'.

description

Specifies a unique description for the list of authentication context classes.

SEE ALSO COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015. All rights reserved.

BIG-IP 2015-11-11 apm saml auth-context-class-list(1)

apm session

NAME

Session - Shows apm session information including session id and all keys and values such as client ip, user name etc

MODULE

apm

SYNTAX

Shows session information with the syntax shown in the following sections.

LIST/DELETE

Lists all the apm session information, session id and all keys and values related to the session. This command can show all the session details or only for some sessions. Deletes all the apm sessions or some sessions at once.

LIST

list session [session_id] [session_id]

options:

all
all-properties
save-to-file

DISPLAY

delete session [session_id] [session_id]

options:

all

EXAMPLES

list session [all-properties] [session_id] [session_id]

Lists details related to the sessions or the given optional session_id.

delete session [all] [session_id] [session_id]

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013, 2015. All rights reserved.

BIG-IP 2015-10-19 apm session(1)

apm sso basic

NAME

basic - Configures a single sign-on HTTP basic authentication configuration object.

MODULE

apm sso

SYNTAX

Configure the basic component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

```
create basic [name]
modify basic [name]
options:
  apm-log-config [[string] | none]
  app-service [[string] | none]
  headers [add | delete | modify | | replace-all-with] {
  location-specific [true | false]
  [name] {
    options:
      app-service [[string] | none]
  hname [[URL] | none]
  hvalue [[integer] | none]
  }
}
password-source [session.sso.token.last.password | none]
username-conversion [enabled | disabled]
username-source [session.sso.token.last.username | none]
```

```
edit basic [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list basic
list basic [ [ [name] | [glob] | [regex] ] ... ]
show running-config basic
show running-config basic [ [ [name] | [glob] |
  [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show basic
show basic [name]
```

DELETE

```
delete basic [name]
```

DESCRIPTION

You can use the basic component to create, modify, display, or delete an SSO HTTP basic authentication configuration object.

EXAMPLES

```
create basic mybasic
```

Creates an SSO basic configuration object named mybasic.

OPTIONS

apm-log-config
Specifies log-setting object to associate with this sso. If this value is empty, logging framework uses log-setting configuration associated with the access profile where sso is used.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

headers

Specifies the name and value of the HTTP header content to be inserted in an HTTP Request that passes through the APM SSO module. The default is none. The options are:

app-service

Specifies the name of the application service to which the HTTP header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the HTTP header. Only the application service can modify or delete the HTTP header.

hname

The name of the HTTP header.

hvalue

The value of the HTTP header.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies a name for the SSO configuration. This option is required.

partition

Displays the partition in which the object resides.

oam-server

Specifies the name of your Oracle Access Manager server. The default value is none.

password source

Specifies the source from which you want SSO to retrieve the password to use to authenticate applications.

username-conversion

Enables or disables conversion of PREWIN2k/UPN username input format to the format for SSO to use. The default value is disabled.

username-source

Specifies the source from which you want SSO to retrieve the username to use to authenticate applications.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2016-03-02 apm sso basic(1)

apm sso form-based

NAME

form-based - Configures a single sign-on form-based configuration object.

MODULE

apm sso

SYNTAX

Configure the form-based component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

create form-based [name]

modify form-based [name]

options:

apm-log-config [[string] | none]

app-service [[string] | none]

external-access-management [oam | none]

form-action [[URL] | none]

form-field [string]

form-method [get | post]

form-password [string]

form-username [string]

headers [add | delete | modify | | replace-all-with] {
[name] {

options:

app-service [[string] | none]

hname [[URL] | none]

hvalue [[integer] | none]

}

}

password-source [session.sso.token.last.password | none]

start-uri [[URLs] | none]

success-match-type [cookie | none | url]

success-match-value [string]

username-source [session.sso.token.last.username | none]

edit form-based [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list form-based

list form-based [[[name] | [glob] | [regex]] ...]

show running-config form-based

show running-config form-based [[[name] | [glob] |
[regex]] ...]

options:

all-properties

non-default-properties

one-line
partition

show form-based
show form-based [name]

DELETE
delete form-based [name]

DESCRIPTION

You can use the form-based component to configure an SSO form-based configuration object.

EXAMPLES

```
create form-based fb_2011_sso { start-uri
"/fb/auth/logon.aspx?url=https://exch2011.mv1.fp.com/fp/&reason=0" form-action "/fp/auth/fpauth.dll" form-
username "username" form-password "password" form-field "destination https://exch2011.mv1.fp.com/fp/"
}
```

Creates an SSO form-based configuration object named fb_2011_sso.

OPTIONS

apm-log-config
Specifies log-setting object to associate with this sso. If this value is empty, logging framework uses log-setting configuration associated with the access profile where sso is used.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

external-access-management
form-action
Specifies the form action URL that is used for HTTP form-based authentication. This is optional. If you do not specify a form action, then Access Policy Manager uses the URI from the request to perform HTTP form-based authentication. The default is none.

form-field
Specifies the hidden form parameters that are required by the authentication server logon form at your location. The default is none. Specify a parameter name, a space, and the parameter value, if any. Multiple parameters can be configured with each "name value" pair in one line. Use edit to add multiple parameters. Please note that create and modify do not allow using new line on the terminal.

form-method
Specifies the form method to use for form-based HTTP authentication. The value is either get or post. The default is post.

If you specify get, Access Policy Manager forces the authentication using HTTP GET rather than authenticating using form-based POST.

form-password
Specifies the parameter names used by the form that is sent the POST request.

form-username
Specifies the parameter names used by the form that is sent the POST request.

headers
Specifies the name and value of the HTTP header content to be inserted in an HTTP Request that passes through the APM SSO module. The default is none.

The options are:

app-service
Specifies the name of the application service to which the HTTP header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the HTTP header. Only the application service can modify or delete the HTTP header.

hname
Specifies the name of the HTTP header.

hvalue
Specifies the value of the HTTP header.

[name]
Specifies a name for the component.

password-source
Specifies the password you want cached for single sign-on.

The default is session.sso.token.last.password.

start-uri
Specifies a URL resource. For example, for FB, it would be /fb/auth/logon.aspx*. For Citrix, /Citrix/XenApp/auth/logon.aspx. This resource must respond with a challenge to a non-authenticated request.

The default is none.

success-match-type

Specifies the method your authentication server uses. If you specify a value for this option, you must also specify a value for success-match-value. The default is none. The options are:

url One or more URLs. The system supports only the wildcard character (*).

cookie

A cookie name.

success-match-value

Specifies the value used to specify either the URL(s) or cookie for the success-match-type option. The default is none.

username-source

Specify the username you want cached for single sign-on. The default is session.sso.token.last.username.

SEE ALSO

basic, kerberos, ntlmv1, ntlmv2

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

BIG-IP 2016-03-02 apm sso form-based(1)

apm sso form-basedv2

NAME

form-basedv2 - Configures a single sign-on form-basedv2 configuration object.

MODULE

apm sso

SYNTAX

Configure the form-basedv2 component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

create form-basedv2 [name]

options:

app-service [[string] | none]
forms [add | replace-all-with] {
 [name] {

 request-value [URIs]

controls [add | replace-all-with] {

 [name] {
 value [string]

 }

 }

 }

 }

modify form-basedv2 [name]

options:

apm-log-config [[string] | none]

app-service [[string] | none]

forms [add | delete | modify | replace-all-with] {
 [name] {

options:

app-service [[string] | none]

attribute-value [[string] | none]

controls [add | delete | modify | replace-all-with] {

 [name] {

 options:

app-service [[string] | none]

secure [true | false]

value [string]

 }

 }

description [[string] | none]

form-order [integer]

id-type [action | id | inputs | name | order]

request-method [get | post]

request-name [[string] | none]

request-negative [true | false]

request-prefix [true | false]

```

request-type [cookie | header | uri]
request-value [[string] | none]
submit-autodetect [true | false]
submit-javascript [[string] | none]
submit-javascript-type [auto | custom | extra]
submit-method post
submit-name [[string] | none]
submit-negative [true | false]
submit-prefix [true | false]
submit-type [cookie | header | uri]
submit-value [[string] | none]
success-match-type [cookie | none | url]
success-match-value [[string] | none]
}
}
headers [add | delete | modify | none | replace-all-with] {
  [name] {
options:
app-service [[string] | none]
description [[string] | none]
name [string]
value [string]
}
}
log-level [alert | crit | debug | emerg | err | info | notice | warn]

```

```

edit form-basedv2 [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

```

reset-stats
reset-stats [ [ [name] | [glob] | [regex] ] ... ]

```

```

DISPLAY
list form-basedv2
list form-basedv2 [ [ [name] | [glob] | [regex] ] ... ]
show running-config form-basedv2
show running-config form-basedv2 [ [ [name] | [glob] |
  [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

```

```

show form-basedv2
show form-basedv2 [name]

```

```

DELETE
delete form-basedv2 [name]

```

DESCRIPTION

You can use the form-basedv2 component to configure an SSO form-basedv2 configuration object. When creating a new SSO form-based v2 configuration object, you must add at least one forms item and within it at least one controls item. You must also provide a value for the request-value option in the forms item.

The SSOv2 module identifies and processes two types of application HTTP requests - logon page requests and credentials submit requests. Logon page requests are identified using the request- set of options. Credentials submit requests, in most cases, are identified automatically. When this fails, you can set the submit-autodetect option to false and use the submit- set of options to identify these requests.

When the SSOv2 module identifies a logon page request, it scans the response trying to find the logon form. If the logon form is found, SSOv2 inserts a JavaScript code that will cause the logon form to be submitted automatically by the browser. The client must support JavaScript.

When the SSOv2 module identifies a credentials submit request, it compares POST data parameter names with form controls defined in the configuration. For a POST data parameter name that has a corresponding form control, the SSOv2 module replaces its value with the control value from the configuration. Control values are usually supplied through session variables, such as session.sso.token.last.username and session.sso.token.last.password. POST data parameters that have no corresponding controls in the configuration are not changed.

The majority of web applications have a single logon page with one logon form. You will need to define a single forms item for these applications. In rare cases when an application has multiple logon pages with different logon forms, you will need to create multiple forms items, one for each logon page/form. If multiple logon pages use the same form, you will need only one forms item with a list of URIs for all logon pages.

Every forms item must include at least one controls item, and can include up to 32 controls items. Each controls item represents an input element of an HTML logon form, such as form fields for entering username and password, and, optionally, any hidden form parameters. The name of the controls item must match the name attribute of the corresponding input tag of the form. For example, if the form has the following HTML tag for entering the username:

the forms item must include a controls item with the name Bugzilla_login. The controls item used for entering

the user's password must have the secure option set to true. The value of a control item should usually be the name of a session variable, starting with the percent (%) sign and enclosed in curly braces ({}); for example, the value for the username control item: `%{session.sso.token.last.username}`. The value can also be a string, or a combination of strings and session variable names.

EXAMPLES

```
create form-basedv2 fbssov2-owa2010 { forms add { owa2010 { controls add { password { secure true value
%\{session.sso.token.last.password\} } username { value %\{session.sso.token.last.username\} } } request-
value /owa/auth/logon.aspx\?replaceCurrent=1 submit-javascript clkLgn() submit-javascript-type extra success-
match-type cookie success-match-value sessionid } } }
```

Creates an SSO form-basedv2 configuration object named fbssov2-owa2010.

```
delete fbsso-owa2010
```

Deletes an SSO form-basedv2 configuration object named fbssov2-owa2010.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

forms

Specifies one or more items, each defining SSO processing of a separate application logon form.

[name]

Specifies the name of the form item. It does not have to match the actual name of the HTML form and can be arbitrary.

The options are:

app-service

Specifies the name of the application service to which the form item belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the form item. Only the application service can modify or delete the form item.

attribute-value

Specifies the value of the HTML tag attribute used to identify the logon form. The attribute could be id, name, or action, and is specified by the id-type option. For other values of the id-type option, this is not used and should be set to none.

controls

Specifies one or more form control items (up to 32) that you want to be processed by SSOv2.

[name]

Specifies the name of the HTML form control item. It must match the name attribute value of the HTML form's input tag.

The options are:

app-service

Specifies the name of the application service to which the control item belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the control item. Only the application service can modify or delete the control item.

secure

Specifies whether the control item represents the HTML input tag of type "password". The default is false.

value

Specifies the value of the control item. This is usually the name of a session variable. If the session variable is not found when the SSO request is processed, the value of the corresponding POST parameter will be empty. The value could also be a literal string or a combination of strings and session variable names.

description

User-defined description.

form-order

Specifies the order of the HTML logon form on the logon page when the id-type option is set to order. Starts with 1.

id-type

Specifies how the HTML logon form is found in the HTML body of the logon page. If there is more than one form on the logon page matching the criteria, the first match is used. The default is inputs.

The options are:

action

The logon form is identified by the value of the tag in the action attribute. The value is specified in the attribute-value option.

id The logon form is identified by the id attribute's value of the tag. The value is specified

in the attribute-value option.

name The logon form is identified by the name attribute's value of the tag. The value is specified in the attribute-value option.

order

The logon form is identified by its relative order on the logon page (starting from 1). The order is specified in the form-order option.

inputs

The logon form is identified by a combination of controls items. The controls in the configuration must have corresponding elements in the form.

request-method

Specifies the HTTP method of the application's request returning logon page. Default is get.

request-name

Specifies the name of the HTTP cookie or the name of the HTTP header used to identify application's request for logon page. The cookie or header is selected by the request-type option. The value of the cookie or header is specified by the request-value option. When the request-type option is set to uri, this option is not used and should be set to none.

request-negative

When set to true, the application's request for logon page will be identified by the absence of the specified cookie or header, or by a failed match against the list of specified URIs. The default is false.

request-prefix

Specifies how the value of the request-value option will be used to match one of the HTTP request cookie, header, or URI. The default is true and specifies a partial match; false specifies an exact match.

request-type

Specifies which element of the HTTP request headers is used to identify the application's request for logon page. The default is uri.

The options are:

cookie

The request is identified by the presence (or absence) of a cookie. The name and value of the cookie are specified by the request-name and request-value options.

header

The request is identified by the presence (or absence) of the HTTP header. The name and value of the header are specified by the request-name and request-value options.

uri The request is identified by a successful (or failed) match against a list of URIs specified by the request-value option, and the request-name option is not used.

request-value

Specifies the value of the HTTP request element that must be matched to identify the request as the application's request for the logon page. This is one of: the cookie value, the header value, or a list of URIs (one per line) as specified by the request-type option. Cookie or header value could be set to none, in which case only the presence of the named cookie or header is checked and the value is not checked. When checking for URI, the value must be specified.

submit-autodetect

When set to true, the application's HTTP request that submits the user's credentials will be identified automatically and other submit- options should not be used. When false, the form submit will be identified using a combination of other submit- options. The default is true.

submit-javascript

Specifies user-provided JavaScript code to be inserted into the logon page to perform automatic form submission when the submit-javascript-type option is set to custom. The custom JavaScript code replaces the code automatically generated by the SSOv2 module. When the submit-javascript-type option is set to extra, it specifies the application's JavaScript functions to call from the automatically generated JavaScript code prior to submitting a logon form. When the submit-javascript-type option is set to auto, this option should be set to none.

submit-javascript-type

Specifies the type of JavaScript code to be inserted into the logon page by the SSOv2 module to perform automatic logon form submission.

The options are:

auto JavaScript code is automatically generated by the SSOv2 module.

custom

JavaScript code is provided by the user in the submit-javascript option.

extra

JavaScript code is automatically generated by the SSOv2 module, and additional JavaScript code provided by the user in the submit-javascript option is inserted before the form submit statement.

submit-method

Specifies the HTTP method of credentials submit request for the application. This must be set to post. This option is not used when submit-autodetect is true.

submit-name

Specifies the name of the HTTP cookie or the name of HTTP header used to identify credentials submit request for the application. The cookie or header is selected by the submit-type option. The value of the cookie or header is specified by the submit-value option. When the submit-type option is set to uri, this option is not used and should be set to none. This option is not used when submit-autodetect is true.

submit-negative

When set to true, the credentials submit request for the application is identified by the absence of a specified cookie or header, or by a failed match against the list of specified URIs. The default is false. This option is not used when submit-autodetect is true.

submit-prefix

Specifies how the value of the submit-value option will be used to match the HTTP request cookie, header, or URI. The default is true and specifies partial match; false specifies exact match. This option is not used when submit-autodetect is true.

submit-type

Specifies which element of HTTP request headers is used to identify the credentials submit request for the application. The default is uri. This option is not used when submit-autodetect is true.

The options are:

cookie

The request is identified by the presence (or absence) of a cookie. The name and value of the cookie are specified by the submit-name and submit-value options.

header

The request is identified by the presence (or absence) of the HTTP header. The name and value of the header are specified by the submit-name and submit-value options.

uri The request is identified by a successful (or failed) match against a list of URIs specified by the submit-value option and the submit-name option is not used.

submit-value

Specifies the value of the HTTP request element that must be matched to identify the request as a credentials submit request for the application. This is one of: the cookie value, the header value, or a list of URIs (one per line) as specified by the submit-type option. Cookie or header value could be set to none, in which case only the presence of the named cookie or header is checked and the value is not checked. When checking for URI, the value must be specified. This option is not used when submit-autodetect is true.

success-match-type

Specifies how the SSOv2 module detects whether the credentials submit request was successful. When the SSOv2 module detects that the credentials submission failed, the SSOv2 configuration used for this HTTP transaction is disabled for the user session. If you specify a value for this option, you must also specify a value for success-match-value. The default is none. The options are:

url Credentials submission was successful if the response contains the HTTP Location header with a value matching one of the URLs specified by the success-match-value option.

cookie

Credentials submission was successful if the response contains the HTTP cookie with the name specified by the success-match-value option.

none No check is performed. If SSO logon fails and the application server redirects back to the logon page that matches the criteria of the logon page request, SSO will be retried, possibly causing authentication loop.

success-match-value

Specifies the value used to detect the success or failure of the SSO logon. When the success-match-type option is set to url, this is a list of URLs. Each URL in the list can contain a single wildcard character (*). When the success-match-type option is set to cookie, this option specifies the name of the cookie. The default is none.

headers

Specifies the name and value of the HTTP header to be inserted in an HTTP Request that passes through the APM SSOv2 module.

[name]

Specifies the name of the headers item.

The options are:

app-service

Specifies the name of the application service to which the HTTP header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the HTTP header. Only the application service can modify or delete the HTTP header.

name Specifies the name of the HTTP header.

value

Specifies the value of the HTTP header.

apm-log-config

Specifies log-setting object to associate with this sso. If this value is empty, logging framework uses log-setting configuration associated with the access profile where sso is used.

log-level

log-level is deprecated. Instead use apm-log-config to customize log-setting.

SEE ALSO

basic, kerberos, ntlmv1, ntlmv2, form-based

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2016-09-28 apm sso form-basedv2(1)

apm sso kerberos

NAME

kerberos - Configures a Kerberos configuration object.

MODULE

apm sso

SYNTAX

Configure the kerberos component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

create kerberos [name]

modify kerberos [name]

options:

account-name [string]

account-password [string]

apm-log-config [[string] | none]

app-service [[string] | none]

headers [add | delete | modify | replace-all-with] {
[name] {

options:

app-service [[string] | none]

hname [[string] | none]

hvalue [[integer] | none]

}

kdc [[string] | none]

location-specific [true | false]

realm [string]

send-authorization [401 | always]

spn-pattern [[string] | none]

ticket-lifetime [[integer] | none]

upn-support [enabled | disabled]

user-realm-source [string]

username-source [string]

edit kerberos [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list kerberos

list kerberos [[[name] | [glob] | [regex]] ...]

show running-config kerberos

show running-config kerberos [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show kerberos

show kerberos [name]

DELETE

delete kerberos [name]

DESCRIPTION

You can use the kerberos component to configure an SSO Kerberos configuration object. Kerberos is an authentication protocol, where both the user and the server verify the other's identity.

EXAMPLES

```
create mykerberos { realm MYREALM.COM account-name apmacount account-password **** }
```

Creates an SSO kerberos configuration object named mykerberos for the realm myrealm.com, where the account name is apmacount and the password is ****.

OPTIONS

account-name

Specifies the name of the Active Directory account configured for delegation. This account must be configured in the server's Kerberos realm (AD Domain). If servers are from multiple realms, each realm (AD Domain) must have its own delegation account. This option is required.

account-password

Specifies the password for the delegation account specified in account-name. This option is required.

apm-log-config

Specifies log-setting object to associate with this sso. If this value is empty, logging framework uses log-setting configuration associated with the access profile where sso is used.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

headers

Specifies custom HTTP headers to insert into a request. The default value is none. The options are:

app-service

Specifies the name of the application service to which the header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the header. Only the application service can modify or delete the header.

hname

Specifies the name of a header to add to a request.

hvalue

Specifies the value of a header to add to a request.

kdc Specifies the IP Address or host name of the Kerberos Key Distribution Center (KDC) for the server's realm. This is normally an Active Directory domain controller. If you leave this empty, the KDC must be discoverable through DNS, for example, BIG-IP system must be able to fetch SRV records for the server realm's domain. If the server realm's domain name is different from the server's realm name, you must specify the server realm's domain name in the /etc/krb5.conf file. Kerberos SSO processing is fastest when KDC is specified by its IP address, slower when specified by host name, and even slower (due to additional DNS queries) when left empty. When a user's realm is different from server's realm, the KDC value must be empty. This is true in cases of cross-realm SSO. The default is none.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies the name for the SSO Kerberos configuration object. This option is required.

realm

Specifies the realm of application server(s), for example, pool members or portal access resource hosts. If the servers are located in multiple realms, each realm requires a separate SSO configuration. You must specify the realm in uppercase letters. The user's realm can be specified through the session.logon.last.domain session variable, and if this variable is not set, then the user's realm is assumed to be the same as the server's realm. This option is required.

send-authorization

Specifies when to submit a Kerberos ticket to the application server(s). The ticket is submitted in an HTTP Authorization header. The header value starts with the word Negotiate, followed by one space and a base64-encoded GSSAPI token containing the Kerberos ticket. If a request contains an Authorization header from the user's browser, it is deleted. The default is always. The options are:

401 The BIG-IP system first forwards the user's HTTP request to the web server without inserting a new Authorization header; however, the browser's Authorization header is deleted. If the server requests authentication by responding with a 401 status code, BIG-IP retries the request with the Authorization header. The Kerberos ticket GSSAPI representation uses the SPNEGO mechanism type (OID 1.3.6.1.5.5.2).

Specifying 401 results in additional BIG-IP/server request round trips in case authentication is required for the request.

always

The BIG-IP system inserts an Authorization header, including the Kerberos ticket, into every HTTP request, whether the request requires authentication or not. The Kerberos ticket GSSAPI representation uses the KRB5 Kerberos 5 mechanism type (OID 1.2.840.113554.1.2.2).

Specifying Always results in the additional overhead of generating a Kerberos token for every request. This is the default value.

spn-pattern

Specifies how the Service Principal Name (SPN) for the server is constructed. For example, HTTP/%s@[server realm name configured in the realm option], where %s will be substituted with the hostname of your server discovered through reverse DNS lookup using the server IP address. Only specify this option when you need non-standard SPN format. The default is none.

ticket-lifetime

Specifies the lifetime of Kerberos tickets obtained for the user. The value represents the maximum ticket lifetime. The actual ticket lifetime may be less by up to 1 hour, because a user's ticket lifetime is the same as the Kerberos Ticket Granting Ticket (TGT) lifetime. A TGT is obtained for the delegation account specified in this configuration. A new TGT is fetched every time the current TGT is older than one hour. The new TGT can only be fetched when an SSO request is processed.

The minimum ticket lifetime is 10 minutes. There is no maximum, however, the ticket lifetime of most AD domains is 10 hours (600 minutes). F5 Networks recommends that you set the ticket lifetime in an SSO configuration above what is specified in an AD domain. The default is 600 minutes.

upn-support

Enables or disables UPN suffix support for Kerberos SSO when integrating into Microsoft Active Directory infrastructure. The default is disabled.

user-realm-source

Session variable name from which Kerberos SSO should read the user's realm. The default is session.logon.last.domain.

username-source

Session variable name from which Kerberos SSO should read the username. The default is session.sso.token.last.username.

SEE ALSO

basic, form-based,ntlmv1, ntlmv2

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2016-09-15 apm sso kerberos(1)

apm sso ntlmv1

NAME

ntlmv1 - Configures a single sign-on (SSO) NT LAN Manager, version 1 (ntlmv1) configuration object.

MODULE

apm sso

SYNTAX

Configure the ntlmv1 component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

create ntlmv1 [name]

modify ntlmv1 [name]

options:

apm-log-config [[string] | none]

app-service [[string] | none]

domain-source [session.logon.last.domain | none]

headers [add | delete | modify | replace-all-with] {
[name] {

options:

app-service [[string] | none]

hname [[string] | none]

hvalue [[integer] | none]

}

location-specific [true | false]

ntlm-domain [[string] | none]

password-source [session.sso.token.last.password | none]

username-conversion [enabled | disabled]

username-source [session.sso.token.last.username | none]

edit ntlmv1 [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ntlmv1

```
list ntlmv1 [ [ [name] | [glob] | [regex] ] ... ]
show running-config ntlmv1
show running-config ntlmv1 [ [ [name] | [glob] |
[regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show ntlmv1
show ntlmv1 [name]
```

```
DELETE
delete ntlmv1 [name]
```

DESCRIPTION

You can use this ntlmv1 component to configure a single sign-on NT LAN Manager, version 1 configuration object.

EXAMPLES

```
create ntlmv1 myntlmv1
Creates an SSO ntlmv1 configuration object named myntlmv1.
```

OPTIONS

apm-log-config
Specifies log-setting object to associate with this sso. If this value is empty, logging framework uses log-setting configuration associated with the access profile where sso is used.

app-service
Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

domain-source
Specifies the Session variable used as a source for the single sign-on user domain. The default is session.logon.last.domain.

headers
Specifies the name and value of the HTTP header content to be inserted in an HTTP Request that passes through the APM SSO module. The default is none.

The options are:

app-service
Specifies the name of the application service to which the HTTP header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the HTTP header. Only the application service can modify or delete the HTTP header.

hname
Specifies the name of the HTTP header.

hvalue
Specifies the value of the HTTP header.

location-specific
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]
Specifies the name for the SSO ntlmv1 configuration object. This option is required.

ntlm-domain
Specifies the static domain setting. If the domain is not retrieved successfully from the source specified in the domain-source option, the system uses this value for the source.

password source
Specifies the source from which you want SSO to retrieve the password to use to authenticate applications. The default is session.sso.token.last.password.

username-conversion
Enables or disables conversion of PREWIN2k/UPN username input format to the format you want to use for SSO. The default is disabled.

username-source
Specifies the source from which you want SSO to retrieve the username used to authenticate applications.

SEE ALSO

basic, form-based,kerberos, ntlmv2

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

apm sso ntlmv2

NAME

ntlmv2 - Configures a single sign-on (SSO) NT LAN Manager, version 2 (ntlmv2) configuration object.

MODULE

apm sso

SYNTAX

Configure the ntlmv2 component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ntlmv2 [name]
modify ntlmv2 [name]
options:
  apm-log-config [[string] | none]
  app-service [[string] | none]
  domain-source [session.logon.last.domain | none]
  headers [add | delete | modify | replace-all-with] {
    [name] {
  options:
    app-service [[string] | none]
    hname [[string] | none]
    hvalue [[integer] | none]
    }
  }
  location-specific [true | false]
  ntlm-domain [[string] | none]
  password-source [session.sso.token.last.password | none]
  username-conversion [enabled | disabled]
  username-source [session.sso.token.last.username | none]
```

```
edit ntlmv2 [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list ntlmv2
list ntlmv2 [ [ [name] | [glob] | [regex] ] ... ]
show running-config ntlmv2
show running-config ntlmv2 [ [ [name] | [glob] |
  [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show ntlmv2
show ntlmv2 [name]
```

DELETE

```
delete ntlmv2 [name]
```

DESCRIPTION

You can use the ntlmv2 component to configure a single sign-on NT LAN Manager, version 2 configuration object.

EXAMPLES

```
create ntlmv2 myntlmv2
Creates an SSO ntlmv2 configuration object named myntlmv2.
```

OPTIONS

apm-log-config
Specifies log-setting object to associate with this sso. If this value is empty, logging framework uses log-setting configuration associated with the access profile where sso is used.

app-service
Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

domain-source
Specifies the Session variable used as a source for the single sign-on user domain. The default is

session.logon.last.domain.

headers

Specifies the name and value of the HTTP header content to be inserted in an HTTP Request that passes through the APM SSO module. The default is none.

The options are:

app-service

Specifies the name of the application service to which the HTTP header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the HTTP header. Only the application service can modify or delete the HTTP header.

hname

Specifies the name of the HTTP header.

hvalue

Specifies the value of the HTTP header.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]

Specifies a name for the sso ntlmv2 configuration object. This option is required.

ntlm-domain

Specifies the static domain setting. If the domain is not retrieved successfully from the source specified in the domain-source option, the system uses this value for the source.

password source

Specifies the source from which you want SSO to retrieve the password to use to authenticate applications. The default is session.sso.token.last.password.

username-conversion

Enables or disables conversion of PREWIN2k/UPN username input format to the format you want to use for SSO. The default is disabled.

username-source

Specifies the source from which you want SSO to retrieve the username used to authenticate applications.

SEE ALSO

basic, form-based,kerberos, ntlmv1

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2016-03-02 apm sso ntlmv2(1)

apm sso oauth-bearer

NAME

oauth-bearer - Configures a single sign-on (SSO) oauth-bearer configuration object.

MODULE

apm sso

SYNTAX

Configure the oauth-bearer component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

create oauth-bearer [name]

modify oauth-bearer [name]

options:

app-service [[string] | none]

headers [add | delete | modify | | replace-all-with] {

location-specific [true | false]

[name] {

options:

app-service [[string] | none]

hname [[URL] | none]

hvalue [[integer] | none]

```

    }
  }
  location-specific [true | false]
  oauth-server [string]

edit oauth-bearer [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

DISPLAY
list oauth-bearer
list oauth-bearer [ [ [name] | [glob] | [regex] ] ... ]
show running-config oauth-bearer
show running-config oauth-bearer [ [ [name] | [glob] |
  [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

show oauth-bearer
show oauth-bearer [name]

```

```

DELETE
delete oauth-bearer [name]

```

DESCRIPTION

You can use the oauth-bearer component to create, modify, display, or delete an SSO oauth-bearer configuration object.

EXAMPLES

```
create oauth-bearer myoauth-bearer
```

Creates an SSO oauth-bearer configuration object named myoauth-bearer.

OPTIONS

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

headers
Specifies the name and value of the HTTP header content to be inserted into an HTTP Request that passes through the APM SSO module. The default is none. The options are:

app-service
Specifies the name of the application service to which the HTTP header belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the HTTP header. Only the application service can modify or delete the HTTP header.

hname
The name of the HTTP header.

hvalue
The value of the HTTP header.

location-specific
Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

[name]
Specifies a name for the SSO oauth-bearer configuration object. This option is required.

partition
Displays the partition in which the object resides.

oauth-server
Specifies the name of your OAuth Server. This option is required.

SEE ALSO

basic, form-based,kerberos, ntlmv1, ntlmv2

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

apm sso saml-resource

NAME

saml-resource - Configures saml resource.

MODULE

apm sso

SYNTAX

Configure a saml-resource using the syntax shown in the following sections.

CREATE/MODIFY

create saml-resource [name]

modify saml-resource [name]

options:

app-service [[string] | none]

customization-group [[string] | none]

description [[string] | none]

location-specific [true | false]

publish-on-webtop [true | false]

sso-config-saml [[string] | none]

edit saml-resource [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list saml-resource

list saml-resource [[[name] | [glob] | [regex]] ...]

show running-config saml-resource

show running-config saml-resource [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete saml-resource [name]

DESCRIPTION

You can use saml-resource component to configure saml resource.

EXAMPLES

```
create saml-resource my_saml_resource { sso-config-saml my_saml_sso_obj publish-on-webtop true }
```

Creates a saml resource named my_saml_resource with saml sso object 'my_saml_sso_obj' and with option to display this resource on full webtop.

```
delete saml-resource my_saml_resource
```

Deletes the saml resource named my_saml_resource.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

customization-group

Specifies the customization group associated with saml resource.

description

Specifies a description for the saml resource. The default is none.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location.

publish-on-webtop

Specifies whether to display the SAML resource on the full webtop or not. Default value is true.

sso-config-saml

Specifies saml sso config object associated with this saml resource. This saml sso object should only have one saml sp connector associated with it.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

apm sso saml-sp-automation

NAME

saml-sp-automation - Specify SAML SP connector automation configuration used to automate creation and management of 'SP Connectors' from the remotely published metadata file(s).

MODULE

apm sso

SYNTAX

Configure the saml-sp-automation component within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

```
create saml-sp-automation [name]
modify saml-sp-automation [name]
options:
  app-service [[string] | none]
  description [[string] | none]
  dns-resolver-name [string]
  frequency [integer]
  metadata-urls [add | delete | modify | none | replace-all-with] {
    name [string] {
      url-value [string]
    }
  }
  serverssl-profile-name [[string] | none]
  sp-obj-name-tag [string]
  sso-config-saml [string]
```

```
edit saml-sp-automation [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list saml-sp-automation
list saml-sp-automation [ [ [name] | [glob] | [regex] ] ... ]
show running-config saml-sp-automation
show running-config saml-sp-automation [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete saml-sp-automation [name]
```

DESCRIPTION

You can use saml-sp-automation to create and manage SAML SP automation objects that are used to create, modify, and delete 'SP Connectors' from the remotely published metadata files.

EXAMPLES

```
create saml-sp-automation my_sp_automation { metadata-urls add { f5 { url-value https://f5.com/metadata.xml }
} dns-resolver-name . sso-config-saml my_saml_idp serverssl-profile-name serverssl } Creates SAML SP
automation object named my_sp_automation bound to a SAML IdP service my_saml_idp with frequency set to 60
minutes with one entry for metadata-url as https://f5.com/metadata.xml, dns-resolver-name as . and serverssl-
profile-name as serverssl.
```

```
list saml-sp-automation
```

Displays a list of SAML SP automation objects.

```
delete saml-sp-automation my_sp_automation
```

Deletes the my_sp_automation SAML SP automation object.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

Specifies the description for the IdP automation object.

dns-resolver-name

Specifies the DNS resolver object to be used for connecting to servers hosting metadata file(s).

frequency

The frequency in minutes at which APM polls the SP metadata files and updates the SP connectors and bindings to the specified SSO SAML server. The default value is 60.

metadata-urls

Specifies a list of one or more URLs containing the metadata files.

serverssl-profile-name

Specifies the SSL profile to be used by the BIG-IP system when connecting to the server hosting metadata file(s).

sp-obj-name-tag

Specifies the name of a tag within the metadata file that contains a value that APM includes in the names of the created SP connectors. If no value is specified, entityID from metadata is used as part of created SP connector name.

sso-config-saml

Specifies the SSO SAML server to which the SP connectors created by this automation are bound.

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2017-10-30 apm sso saml-sp-automation(1)

apm sso saml-sp-connector

NAME

saml-sp-connector - Specify saml sp connector configuration.

MODULE

apm sso

SYNTAX

Configure a saml-sp-connector within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

create saml-sp-connector [name]

modify saml-sp-connector [name]

options:

app-service [[string] | none]

assertion-consumer-services [{

binding [http-artifact | http-post | paos]

index [0 - 65535]

is-default [true | false]

uri [string]

}]

description [[string] | none]

encryption-type [aes128 | aes192 | aes256]

entity-id [string]

import-metadata [string | none]

is-authn-request-signed [true | false]

location-specific [true | false]

metadata-cert [[string] | none]

multi-domain-location [[string] | none]

relay-state [[string] | none]

signature-type [rsa-sha1 | rsa-sha256 | rsa-sha384 | rsa-sha512]

single-logout-binding

single-logout-response-uri [string]

single-logout-uri [string]

sp-certificate [[string] | none]

sp-location [external | internal | internal-multi-domain]

sp-name-qualifier [[string] | none]

want-assertion-encrypted [true | false]

want-assertion-signed [true | false]

want-response-signed [true | false]

edit saml-sp-connector [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

```
list saml-sp-connector
list saml-sp-connector [ [ [name] | [glob] | [regex] ] ... ]
show running-config saml-sp-connector
show running-config saml-sp-connector [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete saml-sp-connector [name]
```

DESCRIPTION

You can use the saml-sp-connector component to create and manage saml sp connectors

EXAMPLES

```
create saml-sp-connector my_saml_sp_connector { entity-id "https://companyx.sp.com/sp" assertion-consumer-
services { { uri "https://companyx.sp.com/acs/" is-default true } } want-assertion-signed true want-response-
signed true want-assertion-encrypted true encryption-type aes256 is-authn-request-signed false sp-certificate
default.crt }
```

Creates a SAML sp-connector named my_saml_sp_connector with security options to encrypt and sign the assertion as well as SAML response.

```
create saml-sp-connector my_saml_sp_connector1 { import-metadata /shared/tmp/sp_metadata.xml}
```

Creates a SAML sp-connector named my_saml_sp_connector1 from metadata file "/shared/tmp/sp_metadata.xml"

```
create saml-sp-connector my_internal_sp_connector { entity-id "https://internal.sp.com" assertion-consumer-
services { { uri "https://internal.sp.com/acs" is-default true } } sp-certificate default.crt sp-location
internal }
```

Creates a SAML sp-connector named my_internal_sp_connector which is load balanced by the same virtual server as this BIG-IP as IdP [identity provider].

```
list saml-sp-connector
```

Displays a list of SAML sp connectors.

```
delete saml-sp-connector my_saml_sp_connector
```

Deletes the my_saml_sp_connector SAML sp connector.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

assertion-consumer-services

List of assertion consumer services (ACS) used by external SP. Each ACS entry contains attributes 'binding', 'index', 'is-default', and 'url'. Each ACS must contain a valid URL, and a unique 'index'. One ACS entry must be set as default.

assertion-consumer-binding

This attribute is DEPRECATED. Use assertion-consumer-services instead.

assertion-consumer-uri

This attribute is DEPRECATED. Use assertion-consumer-services instead.

description

Specifies a unique description for saml sp connector. The default is none.

encryption-type

Specifies the type of encryption BIG-IP as IdP should use to encrypt the assertion. Default is aes128.

entity-id

Specifies a unique ID to identify SP pointed by sp connector.

import-metadata

Specifies the metadata file to be used to create sp connector object. For example: create saml-sp-connector my_saml_sp_connector1 { import-metadata /shared/tmp/sp_metadata.xml}

is-authn-request-signed

Specifies whether SP signs authentication requests while sending them to BIG-IP as IdP. The default value for this is false.

location-specific

Objects of this class might have location specific attribute(s). Admin can indicate if object is location specific by setting it to true.

metadata-cert

Specifies the certificate to be used to verify the signature of metadata imported from a file.

multi-domain-location

Specifies the scheme, hostname, and (optionally) port of the virtual server on this BIG-IP behind which this SP is located, e.g. "https://application.f5.com". This configuration is required only when sp-location attribute is configured as 'internal-multi-domain'

relay-state

Specifies the value sent to the SP by BIG-IP as IdP as part of the response. This value is only used if the SP did not send RelayState as part of the authentication request.

signature-type

Signature algorithms to be used for digital signing of SAML messages. Default value is rsa-sha1.

single-logout-binding

This attribute is reserved for future functionality.

single-logout-response-uri

A URI where this BIG-IP as IdP will send single logout (SLO) responses.

single-logout-uri

A URI where this BIG-IP as IdP will send single logout (SLO) requests.

sp-certificate

Specifies SP certificate used by BIG-IP as IdP to verify the signature of authentication request.

sp-location

Specifies the location of SP from network topology viewpoint. Default value external should be used with SAML WebSSO profile. This value indicates that SP is located externally from BIG-IP perspective, and therefore SP is reachable directly by the user-agent. internal - indicates that configured SP is located behind the virtual server that hosts BIG-IP IdP, and therefore SP is not reachable directly by the client. internal-multi-domain - indicates that BIG-IP is configured for multi-domain SSO, and therefore SP is located behind different virtual server of this BIG-IP.

sp-name-qualifier

Optionally qualifies an identifier with the name of a service provider or affiliation of providers.

want-assertion-encrypted

Specifies whether SP requires encrypted assertions. The default value for this attribute is false

want-assertion-signed

Specifies whether SP requires signed assertions. The default value for this attribute is true

want-response-signed

Specifies whether SP requires signed SAML responses. The default value for this attribute is false

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2018-01-10 apm sso saml-sp-connector(1)

apm sso saml

NAME

saml - Specify SAML SSO configuration.

MODULE

apm sso

SYNTAX

Configure the saml within the sso module using the syntax shown in the following sections.

CREATE/MODIFY

create saml [name]

modify saml [name]

options:

apm-log-config [[string] | none]

app-service [[string] | none]

artifact-resolution-service-name [name | none]

assertion-validity [integer]

attributes [none | {

{

name [[string] | none],

name-format [basic | unspecified | uri],

multi-values {

[string]

},

encrypt [true | false],

encryption-type [aes128 | aes192 | aes256]

}

}]

```

auth-context-method [string | none]
description [[string] | none]
encrypt-subject [true | false]
encryption-type-subject [aes128 | aes192 | aes256]
entity-id [string]
export-metadata [no-signing | with-signing]
idp-certificate [string | none]
idp-certificate-session-var [string | none]
idp-host [string | none]
idp-scheme [http | https]
idp-signkey [string | none]
idp-signkey-session-var [string | none]
key-transport-algorithm [ rsa-oaep | rsa-v1.5 ]
location-specific [true | false]
log-level [alert | crit | debug | emerg | err | info | notice | warn]
metadata-cert [[string] | none]
metadata-file [[string] | none]
metadata-signkey [string | none]
name-qualifier [[string] | none]
saml-profiles [add | delete | modify | none | replace-all-with] {
[ecp | web-browser-sso]
}
sp-connectors [add | delete | modify | none | replace-all-with] {
[string]
}
subject-type [email-address | kerberos | transient | win-domain-qualified-name | entity | persistent | unspecified | x509-subject]
subject-value [ string | none ]

```

edit saml [[[name] | [glob] | [regex]] ...]

```

options:
  all-properties
  non-default-properties

```

DISPLAY

list saml

list saml [[[name] | [glob] | [regex]] ...]

show running-config saml

show running-config saml [[[name] | [glob] | [regex]] ...]

```

options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition

```

DELETE

delete saml [name]

DESCRIPTION

You can use the saml component to create and manage SAML SSO objects.

EXAMPLES

```

create saml my_saml_sso_obj { entity-id "https://myidpvs.big-ip.com/idp" subject-type email-address subject-
value test@mycompany.com idp-certificate default.crt idp-signkey default.key sp-connectors add { google_apps
salesforce }}

```

Creates a SAML SSO object named my_saml_sso_obj with SP connectors "google_apps" and "salesforce"

```

create saml my_saml_sso_obj1 { entity-id "https://myidpvs.big-ip.com/idp" subject-type email-address subject-
value test@mycompany.com idp-certificate default.crt idp-signkey default.key sp-connectors add { google_apps
sp_salesforce } attributes {{name "group" multi-values { "PD" "Admin" }} {name "title" multi-values {
"engineer1" }}} }

```

Creates a SAML SSO object named my_saml_sso_obj1 with attributes "group" and "title".

```
list saml
```

Displays list of SAML SSO objects.

```
delete saml my_saml_sso_obj
```

Deletes the my_saml_sso_obj SAML SSO object.

OPTIONS

```
apm-log-config
```

Specifies log-setting object to associate with this saml. If this value is empty, logging framework uses log-setting configuration associated with the access profile where sso is used.

```
app-service
```

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

```
assertion-validity
```

Specifies assertion validity period in seconds.

```
artifact-resolution-service-name
```

Specifies the artifact resolution service to be used by this BIG-IP as IdP to receive artifacts and resolve them for assertions.

```
attributes
```

Specifies list of attributes as part of assertion. Both attribute name and values can be session variables. Property 'value' is DEPRECATED; "multi-values" must be used instead. "name-format" can be used to optionally specify format of the attribute name.

```
create saml my_saml_sso_obj1 { entity-id "https://myidpvs.big-ip.com/idp" subject-type email-address
subject-value test@mycompany.com idp-certificate default.crt idp-signkey default.key sp-connectors add {
google_apps sp_salesforce } attributes { {name "group" multi-values {
"%{session.ldap.last.attr.primarygroup}"}} {name "name" multi-values { "firstName" "lastName" } name-
format basic}} }
```

Creates a SAML SSO object named my_saml_sso_obj1 with attributes "group" and "name".

auth-context-method

Specifies an authentication context method used by this BIG-IP as IdP when creating assertions. This attribute can be a session variable.

description

Specifies a unique description for SAML SSO object. The default is none.

encrypt-subject

Set to true if assertion 'Subject' must be encrypted. Default value is false.

encryption-type-subject

Encryption algorithm used to encrypt 'Subject' element in assertion. Default value is aes128.

entity-id

Specifies unique identifier for BIG-IP as IdP. Typically, 'entity-id' is a URI that points to the BIG-IP virtual server that is going to act as a SAML IdP. In case 'entity-id' is not a valid URL, the idp-host attribute is required. Examples of valid configuration include "https://mycompany-idp", "idp:my:company", and "idp.my.company.com"

export-metadata

You can simplify SAML configuration using metadata files. When you use APM as an IdP, you can export metadata for IdP. You can save metadata to a file and give it to the SP to enable SP to import SP's SAML configuration or enable SP to use information from the metadata file to configure the IdP. You can choose to sign metadata while exporting it for better security.

For example:

1. Exporting metadata with signing. This requires metadata-signkey and metadata-cert files.

```
modify saml my_saml_sso_obj {export-metadata with-signing metadata-file /shared/idp_signed_metadata.xml metadata-cert default.crt m
```

2. Exporting metadata with no signing.

```
modify saml my_saml_sso_obj {export-metadata no-signing metadata-file /shared/idp_metadata.xml}
```

idp-certificate

BIG-IP includes this certificate in the SAML IdP metadata that you export. After the SAML IdP metadata is imported on the SP, the SP can use this certificate to verify the signature of assertion sent by this BIG-IP as IdP.

idp-certificate-session-var

Specifies the certificate this BIG-IP as IdP will use to sign SAML messages including SAML assertion. This attribute must be specified in session variable format. This attribute is mutually exclusive with

idp-host

Hostname of this BIG-IP as IdP. This attribute is required when "entity-id" is not a valid URL.

idp-scheme

Scheme used by this BIG-IP as IdP. This attribute is only used when idp-host is not empty. Default value is https.

idp-signkey

Specifies the private key used for signing assertion by BIG-IP as IdP.

idp-signkey-session-var

Specifies the signing key this BIG-IP as IdP will use to sign SAML messages including SAML assertion. This attribute must be specified in session variable format. This attribute is mutually exclusive with

key-transport-algorithm

Specifies the key transport algorithm to be used for encrypted attributes, subject-value, or assertion. Default and recommended value is rsa-oaep. rsa-v1.5 is NOT RECOMMENDED due to security risks associated with the algorithm, and should NOT be used except for compatibility with older applications.

location-specific

Objects of this class might have location specific attribute(s). Admin can indicate if object is location specific by setting it to true.

log-level

log-level is deprecated. Instead use apm-log-config to customize log-setting.

metadata-cert

Specifies the certificate with public key of the key pair used in signing the metadata. See export-metadata for more information on metadata export functionality. This is the certificate to include in

signed metadata when we export metadata. This might or might not be IdP certificate.

metadata-file

Specifies the file to which metadata is saved. See export-metadata for more information on metadata export functionality.

metadata-signkey

This specifies the key that is used to sign IdP's metadata. See export-metadata for more information on metadata export functionality.

name-qualifier

Specifies the security or administrative domain of the IdP (this BIG-IP system). This value usually matches IdP Entity ID.

saml-profiles

List of SAML profiles enabled on this BIG-IP as IdP. Default value is web-browser-sso.

sp-connectors

Specifies list of SP connectors associated with this SAML SSO object. When this SSO object is assigned to SAML resource then only one entry is allowed for SP connectors. If SAML SSO object is assigned to access profile then you can add multiple SAML SP connectors.

subject-type

Specifies type of the subject to be used while creating SAML assertion.

subject-value

Specifies the value of the subject to be included inside SAML assertion. This can be a session variable. For example: `%{session.last.logonname}`, `%{session.ad.last.attr.userEmail}`

SEE ALSO

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2017-11-29 apm sso saml(1)

apm swg-content-type

NAME

swg-content-type - Containers for the SWG content types

MODULE

apm

SYNTAX

List the swg-content-type objects. Creating, modifying, or deleting is not allowed.

Here are the valid values: All-Compressed, All-Executable, All-Images, Application-flash, Text-html, Text-pdf

DISPLAY

list swg-content-type

list swg-content-type [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DESCRIPTION

Indicates the MIME type of the response content. Used in SWG Response Analytics.

EXAMPLES

list apm swg-content-type

Lists all swg-content-type objects.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2017. All rights reserved.

BIG-IP 2017-09-05 apm swg-content-type(1)

apm swg-scheme

NAME

swg-scheme - Configures an SWG Scheme (The swg-scheme object is deprecated)

MODULE

apm

SYNTAX

Configure a swg-scheme component within the apm module using the syntax shown in the following sections.

CREATE/MODIFY

Each swg-scheme consists of the following.

```
create swg-scheme [name]
options:
  app-service [[string] | none]
  description [[string] | none]
```

```
modify swg-scheme [name]
options:
  app-service [[string] | none]
  description [[string] | none]
```

DISPLAY

```
list swg-scheme
list swg-scheme [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DESCRIPTION

Configures an swg-scheme object. Note that the swg-scheme object is deprecated as of v15.1.0.

EXAMPLES

```
create apm swg-scheme scheme1 { app-service none description "My SWG Scheme" }
```

Creates a new swg-scheme.

```
modify swg-scheme scheme1 { description "My SWG Scheme" }
```

Modify a swg-scheme by modifying its description string.

OPTIONS

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description

Specifies a user-defined description for the swg scheme.

SEE ALSO

apm url-filter

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2019-06-24 apm swg-scheme(1)

apm url-filter

NAME

url-filter - Configures URL filters for URL classification and filtering

MODULE

apm

SYNTAX

Configure a url-filter component within the apm module using the syntax shown in the following sections.

CREATE/MODIFY

Each url-filter consists of two url-category lists: a list of allowed URL categories and a list of blocked URL categories. The requests for URLs contained in the allowed list are allowed to pass unfettered, whereas requests for URLs in the blocked list will not go out into the Internet.

```
create url-filter [name]
options:
  allowed-categories [add | delete | modify | replace-all-with] {
[string]
}
  blocked-categories [add | delete | modify | replace-all-with] {
[string]
}
```

```
modify url-filter [name]
  allowed-categories [add | delete | modify | replace-all-with] {
[string]
}
  blocked-categories [add | delete | modify | replace-all-with] {
[string]
}
```

DISPLAY

list url-filter

```
list url-filter [ [ [name] | [glob] | [regex] ] ... ]
```

options:

- all-properties
- non-default-properties
- one-line
- partition

CP

```
cp url-filter [source-name] [target-name]
```

DESCRIPTION

Configures a url-filter

NOTE: A url-filter can have a large number of URL categories in each list. To facilitate the creation of url-filter, you can create a new url-filter by copying from an existing url-filter. Then modify each list by adding or removing url-categories to suit your needs.

EXAMPLES

```
create url-filter my-url-filter allowed-categories add { Business_and_Economy Education } blocked-categories
add { Adult_Content Shopping }
```

Creates a new url-filter.

```
modify url-filter my-own-filter allowed-categories delete { Education }
```

Modify a url-filter by deleting a URL category from the allowed list.

```
cp url-filter existing-filter another-filter
```

Create a new url-filter by copying from an existing filter.

OPTIONS

allowed-categories

Specifies the URL categories that should be allowed to pass.

description

Specifies a unique description for the URL filter.

blocked-categories

Specifies the URL categories that should be blocked.

SEE ALSO

sys url-db download-result sys url-db download-schedule and sys url-db url-category

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013, 2015. All rights reserved.

BIG-IP 2015-07-22 apm url-filter(1)

asm device-sync

NAME

device-sync - Contains the ASM timestamp for each device in the group.

MODULE

asm

SYNTAX

Retrieve the list of the device-sync values using the syntax shown in the following section.

DISPLAY

```
list device-sync
```

```
list device-sync [ [ [name] | [glob] | [regex] ] ... ]
```

DESCRIPTION

Use this command to display the current values of the device-sync object, i.e. ASM change times for all devices in the group. This object is designed for internal purposes only (incremented on every ASM change), so do not try to create, modify, or delete it manually.

EXAMPLES

```
list device-sync
```

Displays all last ASM change times of the device group.

SEE ALSO

tmssh, list, glob, regex

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2011. All rights reserved.

BIG-IP 2012-03-02 asm device-sync(1)

asm http-method

NAME

http-method - Lists the available HTTP request methods that can be used in the context of the Application Security Manager(TM).

MODULE

asm

SYNTAX

Retrieve the list of the http-method values using the syntax shown in the following sections.

DISPLAY

```
list http-method
```

```
list http-method [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all
app-service
default-act-as
one-line
partition
recursive
```

DESCRIPTION

Use this command to display the possible values of the http-method object to be used in the context of the Application Security Manager. These possible values include predefined and user-defined allowed methods for all security policies, and also are intended to be used in filters of Application Security Logging and in HTTP security profiles.

EXAMPLES

```
list http-method
```

Displays all the HTTP methods supported by the ASM.

OPTIONS

```
app-service
```

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

default-act-as

Displays the HTTP request method, either GET or POST, based on how you have instructed the system to treat the listed method name; a predefined method has its system default and a user-defined allowed method is configured in the security policy.

partition

Displays the administrative partition within which the component resides.

SEE ALSO

glob, list, regex, security http profile, security log profile, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-06-12 asm http-method(1)

asm httpclass-asm

NAME

httpclass-asm - configure initial ASM settings for applications. This component has been deprecated as of BIG-IP v11.3.0, please use the policy component in the asm module instead.

MODULE

asm

SYNTAX

Configure the httpclass-asm component within the asm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create httpclass-asm [name]
```

```
modify httpclass-asm [name]
```

options:

```
active-policy-name [string]
```

```
app-service [[string] | none]
```

```
language [language]
```

```
predefined-policy [predefined-policy]
```

DISPLAY

```
list httpclass-asm
```

```
list httpclass-asm [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config httpclass-asm
```

```
show running-config httpclass-asm [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
partition
```

DELETE

```
delete httpclass-asm [name]
```

DESCRIPTION

Use this command to create, modify, display, or delete an httpclass-asm profile that configures ASM security policies. Changing/setting attributes for an httpclass-asm profile affects the ASM security policy with the same name. Note that modifying the language of an existing profile reconfigures the ASM security policy and deletes the configurations, log entries and statistics of the security policy. This is for advanced usage - this command is intended to be used by the application templates system (iApps(tm)).

EXAMPLES

```
create asm httpclass-asm my_class active-policy-name my_class_policy language utf-8 predefined-policy  
POLICY_TEMPLATE_RAPID_DEPLOYMENT_HTTP
```

Creates a custom httpclass-asm profile named my_class that causes ASM to configure a security policy that uses the utf-8 application language and the Rapid Deployment security policy.

```
list httpclass-asm
```

Displays the properties of all httpclass-asm profiles.

OPTIONS

active-policy-name

Specifies the name of the active security policy. This property has been deprecated. As of BIG-IP v11.1.0, the active security policy name is identical to the HTTP class profile's name.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

language

Specifies the language of the web application that the ASM security policy is protecting. Use autocomplete or list /asm webapp-language to get the list of supported languages.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

predefined-policy

Specifies a predefined security policy for a web application. This security policy was prebuilt to provide out of the box security for a known application. Use autocomplete to get a list of applications for which ASM has predefined policies.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-04-12 asm httpclass-asm(1)

asm policy

NAME

policy - Configures an application security policy.

MODULE

asm

SYNTAX

Configure the policy component within the asm module using the syntax shown in the following sections.

CREATE

create policy [name]

options:

[active | inactive]
app-service [[string] | none]
blocking-mode [enabled | disabled]
description [[string] | none]
encoding [[name] | none]
policy-builder [enabled | disabled]
policy-template [name]
policy-type [security | parent]
parent-policy [name]

MODIFY

modify policy [name]

options:

[active | inactive]
app-service [[string] | none]
blocking-mode [enabled | disabled]
description [[string] | none]
encoding [[name] | none]
policy-builder [enabled | disabled]
policy-template [name]

DISPLAY

list policy [[[name] | [glob] | [regex]] ...]

show running-config policy [[[name] | [glob] | [regex]] ...]

options:

all-properties
one-line
partition
virtual-servers

DELETE

delete policy [name]

SAVE

save policy [name]

options:

overwrite

bin-file [filename]
min-xml-file [filename]
xml-file [filename]

LOAD

load policy [name]
options:
 overwrite
 file [filename]
 xml-string [string]

PUBLISH

publish policy [name]

DESCRIPTION

You can use the policy component to create, modify, display, delete, save, load, or publish an application security policy for use with Application Security Manager functionality.

Note: To display all policy properties available in tmsh, including initial settings used by iApp and advanced configuration accessible in ASM GUI, specify the all-properties option or the detailed properties. By default, only initial properties are displayed: encoding, policy-template and [active | inactive].

Note: The modify command with the properties encoding and/or policy-template causes ASM to reconfigure the security policy and clear all its former data.

Note: The policy-type cannot be modified after the creation of the policy.

Note: The parent-policy can only, optionally, be set while creating a policy with policy-type set to security.

EXAMPLES

```
create policy my_asm_policy encoding utf-8
```

Creates a new policy named my_asm_policy with the default language encoding, policy-type set to security and no parent-policy.

```
modify policy my_asm_policy active
```

Activates the inactive policy named my_asm_policy.

```
create policy my_parent_asm_policy encoding utf-8 policy-type parent
```

Creates a new policy named my_parent_asm_policy with the default language encoding.

```
create policy my_security_asm_policy policy-type security parent-policy my_parent_asm_policy
```

Creates a new policy named my_security_asm_policy, policy-type set to security and my_parent_asm_policy set as the parent policy.

```
list policy
```

Displays the properties of all application security policies.

```
save policy my_asm_policy xml-file my_asm_policy.xml
```

Exports the policy named my_asm_policy to the XML file /var/tmp/my_asm_policy.xml.

```
load policy my_asm_policy overwrite file /tmp/my_asm_policy.plc
```

Imports the policy named my_asm_policy from the file /tmp/my_asm_policy.plc and overwrites the policy if it already exists.

```
publish policy my_asm_policy
```

Applies the active policy named my_asm_policy.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

[active | inactive]

Activates or deactivates the policy for later association with L7 policies and virtual servers. The default value is inactive.

bin-file

Specifies the exported file name to be saved in binary format when using the save command. The file name should be simple (not a full path); it is saved to the /var/tmp directory on the system.

blocking-mode

Specifies whether the system blocks a request that triggers a security policy violation or only logs the violation event (transparent mode).

description

Specifies an optional description of the security policy.

encoding

Specifies the language encoding, which determines how the security policy processes the character sets. This property corresponds to the language property of the httpclass-asm component.

`file` Specifies the file name from which the policy is going to be imported when using the load command. A full path should be specified.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`min-xml-file`
Specifies the exported file name to be saved in compact XML format when using the save command. The file name should be simple (not a full path); it is saved to the /var/tmp directory on the system. To display the XML output immediately, omit this property, the properties xml-file and bin-file.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, modify, save, and publish. If it is not specified for the load command, the policy name will be taken from the imported settings.

`overwrite`
Specifies that the policy file for the save command or the policy component for the load command can be overwritten if it exists.

`partition`
Displays the administrative partition within which the component resides.

`policy-builder`
Enables or disables automatic policy building.

`policy-template`
Specifies whether the security policy is based on a predefined security policy template, and if so, which one. If you create or modify a security policy based on a template, the system automatically configures the new security policy according to the conditions of the template. This property corresponds to the predefined-policy property of the httpclass-asm component.

`policy-type`
Specifies the security policy type, which cannot be changed after you create the policy. The parent policy type cannot be active and cannot have a parent. The security policy type may or may not have a single parent policy defined.

`parent-policy`
Optionally, specifies the name of an existing policy, of policy-type parent, to be set as the parent policy, while creating a policy-type security.

`regex`
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

Note: This component supports matching by the regex expression only when displaying the initial policy properties.

`virtual-servers`
Displays the name of the protected virtual server, or virtual servers, which have attached to them the security policy via L7 policies.

`xml-file`
Specifies the exported file name to be saved in XML format when using the save command. The file name should be simple (not a full path); it is saved to the /var/tmp directory on the system. To display the XML output immediately, omit this property, the properties min-xml-file and bin-file.

`xml-string`
Specifies the XML document from which the policy is going to be imported when using the load command.

SEE ALSO

asm predefined-policy, asm webapp-language, create, delete, glob, list, load, ltm policy, ltm virtual, modify, publish, regex, save, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2012. All rights reserved.

BIG-IP 2016-09-14 asm policy(1)

asm predefined-policy

NAME

predefined-policy - Lists the available predefined policies that can be used in the context of the httpclass-asm profile.

MODULE
asm

SYNTAX
Retrieve the list of the predefined-policy values using the syntax shown in the following sections.

DISPLAY
list predefined-policy
list predefined-policy [[[name] | [glob] | [regex]] ...]
options:
all
one-line

DESCRIPTION
Use this command to display the possible values of the predefined-policy object to be used in the context of the httpclass-asm profile. This is for advanced usage; this command is intended for use by the application templates system (iApps).

EXAMPLES
list predefined-policy

Displays all the predefined policies supported by the ASM.

SEE ALSO
asm httpclass-asm, glob, list, regex, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2011. All rights reserved.

BIG-IP 2013-01-27 asm predefined-policy(1)

asm response-code

NAME
response-code - Lists the available HTTP response status codes that can be used in the context of the Application Security Manager.

MODULE
asm

SYNTAX
Retrieve the list of the response-code values using the syntax shown in the following sections.

DISPLAY
list response-code
list response-code [[[number] | [glob] | [regex]] ...]
options:
all
app-service
name
one-line

DESCRIPTION
Use this command to display the possible values of the response-code object to be used in the context of the Application Security Manager. These possible values are predefined and intended to be used in filters of Application Security Logging.

EXAMPLES
list response-code

Displays all the response codes supported by the ASM.

OPTIONS
app-service
Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

name Displays a well-known textual meaning of the HTTP response code.

SEE ALSO
glob, list, regex, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2013-01-27 asm response-code(1)

asm webapp-language

NAME

webapp-language - Lists the available languages that can be used in the context of the httpclass-asm profile.

MODULE

asm

SYNTAX

Retrieve the list of the webapp-language values using the syntax shown in the following sections.

DISPLAY

```
list webapp-language
list webapp-language [ [ [name] | [glob] ... ]
options:
  all
  one-line
```

DESCRIPTION

Use this command to display the possible values of the webapp-language object to be used in the context of the httpclass-asm profile. This is for advanced usage - this command is intended to be used by the application templates system.

EXAMPLES

```
list webapp-language
```

Displays all the languages supported by the ASM.

SEE ALSO

asm httpclass-asm, glob, list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2011. All rights reserved.

BIG-IP 2013-01-27 asm webapp-language(1)

auth

auth apm-auth

NAME

apm-auth - Configures an APM-based authentication object for implementing access policy execution-based authentication of BIG-IP(r) system users.

MODULE

auth

SYNTAX

Configure the apm-auth component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

```
create apm-auth [name]
modify apm-auth [name]
options:
  profile-access [string]
```

```
edit apm-auth [ [ [name] | [glob] | [regex] ] ...]
```

options:
all-properties
non-default-properties

DISPLAY

```
list apm-auth  
list apm-auth [ [ [name] | [glob] | [regex] ] ...]  
show running-config apm-auth  
show running-config apm-auth [ [ [name] | [glob] | [regex] ] ...]  
options:  
all-properties  
non-default-properties  
one-line  
partition
```

DELETE

```
delete apm-auth [name]
```

DESCRIPTION

You can configure APM-based authentication to execute an access policy for BIG-IP system users to authenticate and authorize them.

APM authentication methods like HTTP, AD/LDAP, TACACS+ authentication can be used in box authentication. To authenticate BIG-IP system users, to do this, create an access profile, create APM configuration object with the above access profile, and then activate the object.

The following steps describe how to configure APM-based authentication for BIG-IP system users:

1. Use the profile-access component in the apm module to create an access profile.
2. Use the apm-auth component in the auth module to map an access profile to an APM-based authentication object.
3. To activate APM-based authentication for BIG-IP system users, run the command sequence `modify / auth source type apm-auth`

EXAMPLES

```
create apm-auth system-auth {profile-access apm-profile}
```

Creates an APM-based authentication object named system-auth

```
delete apm-auth system-auth
```

Deletes the APM-based authentication object named system-auth.

OPTIONS

description
User-defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition
Displays the administrative partition within which the component resides.

profile-access
Specifies the access profile that the system must use for APM-based authentication. You must specify an access profile when you create an APM-based configuration object.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

auth user, create, delete, glob, list, modify, regex, run, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-06-17 auth apm-auth(1)

NAME

cert-ldap - Configures an LDAP configuration object for implementing Single Sign On based on a valid client certificate for BIG-IP(r) system users. The user is required to properly configure the Certificate Authority so that unique identifying attributes appear in the subjectName or subjectAltName fields of signed client certificates; the OCSP responder so that it is available to the BIG-IP at the time a client certificate is presented; and the LDAP server so that it includes the required attributes from the client certificate and the corresponding user name.

MODULE

auth

SYNTAX

Configure the cert-ldap component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

create cert-ldap [name]

modify cert-ldap [name]

options:

bind-dn [[account dn] | none]

bind-pw [none | [password]]

bind-timeout [integer]

check-host-attr [disabled | enabled]

check-roles-group [disabled | enabled]

debug [disabled | enabled]

description [string]

filter [[filter name] | none]

idle-timeout [integer]

ignore-auth-info-unavail [no | yes]

ignore-unknown-user [disabled | enabled]

login-attribute [[account name] | none]

login-filter [[string] | none]

login-name [[ldap attribute] | none]

port [[name] | [integer]]

scope [base | one | sub]

search-base-dn [[search base dn] | none]

search-timeout [integer]

servers [add | delete | replace-all-with] {

[[ip address] | [server name] ...] }

ssl [disabled | enabled]

ssl-ca-cert-file [[file name] | none]

ssl-check-peer [disabled | enabled]

ssl-ciphers [[string] | none]

ssl-client-cert [[string] | none]

ssl-client-key [[string] | none]

ssl-cname-field [subjectname-cn | san-other | san-email

san-dns | san-x400 | san-dirname | san-ediparty

san-uri | san-ipadd | san-rid]

ssl-cname-otheroid [[OID in dotted-decimal] | none]

sso [on | off]

version [integer]

warnings [disabled | enabled]

edit cert-ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list cert-ldap

list cert-ldap [[[name] | [glob] | [regex]] ...]

show running-config cert-ldap

show running-config cert-ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete cert-ldap [name]

DESCRIPTION

The CERT-LDAP authentication mode is required to provide Single Sign On capability to the control plane based on a valid client certificate. This mode involves configuring an Apache server to initiate a client certificate request, perform certificate validation against an OCSP server, and then authenticate/authorize certificate credentials against a configured remote LDAP server or a Microsoft(r) Windows(r) Active Directory(r). The mode is not based on basic HTTP authentication (that is, user name and password). CERT-LDAP mode is equivalent to LDAP mode with custom attributes.

To authenticate BIG-IP system users when their authentication data is stored on a remote LDAP server, you create an LDAP configuration object, and then activate the object. Make sure that Apache is configured to support the client certificate validation.

To configure CERT-LDAP authentication for BIG-IP system users:

1. Use the cert-ldap component in the auth module to configure an LDAP configuration object.

2. To activate LDAP authentication for BIG-IP system users, run the command sequence modify / auth source type cert-ldap

EXAMPLES

```
create cert-ldap bigip_cert_ldap_auth servers add {my_ldap_server}
```

Creates a configuration object named bigip_cert_ldap_auth.

```
delete cert-ldap bigip_cert_ldap_auth
```

Deletes the configuration object named bigip_cert_ldap_auth.

OPTIONS

bind-dn

Specifies the distinguished name of an account to which to bind to perform searches. This search account is a Read-only account. You can also use the admin account as the search account. If an administrative distinguished name is not specified, then a bind is not attempted. The default value is none.

Note: If the remote server is a Microsoft Windows Active Directory server, the distinguished name must be in the form of an email address.

bind-pw

Specifies the password for the search account created on the LDAP server. This option is required if you enter a value for the bind-dn option. The default value is none.

bind-timeout

Specifies a bind timeout limit, in seconds. The default value is 30.

check-host-attr

Confirms the password for the bind distinguished name. This option is optional. The default value is disabled.

check-roles-group

Specifies whether to verify a user's group membership given in the remote-role definitions, formatted as *member*of="group-dn". The default value is disabled.

debug

Enables or disables syslog-ng debugging information at the LOG DEBUG level. The default value is disabled. F5 Networks does not recommend using this option for normal configuration.

description

User defined description.

filter

Specifies a filter. Use this option for authorizing client traffic. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

group-dn

Specifies the group distinguished name. The system uses this option for authorizing client traffic. The default value is none.

group-member-attribute

Specifies a group member attribute. The system uses this option for authorizing client traffic. The default value is none.

idle-timeout

Specifies the idle timeout, in seconds, for connections. The default value is 3600 seconds.

ignore-auth-info-unavail

Specifies whether the system ignores authentication information if it is not available. The default value is no.

ignore-unknown-user

Specifies whether the system ignores a user that is unknown. The default value is disabled.

login-attribute

Specifies a logon attribute. Normally, the value for this option is uid; however, if the server is a Microsoft Windows Active Directory server, the value must be the account name samaccountname (not case-insensitive). The default value is none.

login-filter

Specifies the filter to be applied on the CN of the client certificate. This filter is a regular expression to extract required information from CN of client certificate which will be used to match against LDAP search results. The default is disabled.

login-name

Specifies the LDAP attribute holding the client name. (The client name is extracted from the client certificate as specified by ssl-cname-field.) The default is disabled.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

port Specifies the port number or name for the LDAP service. Port 389 is typically used for non-SSL and port

636 is used for an SSL-enabled LDAP service. The default value is ldap.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

scope

Specifies the search scope. The default value is sub. The possible values are:

base The search scope is base object. The base value is almost never useful for name service lookups.

one The search scope is one level.

sub The search scope is a subtree.

search-base-dn

Specifies the search base distinguished name. The default value is none.

search-timeout

Specifies the search timeout, in seconds. The default value is 30.

servers

Specifies the LDAP servers that the system must use to obtain authentication information. You must specify a server when you create an LDAP configuration object.

ssl Enables or disables SSL functionality. The default is disabled.

Note that when you use tmsh to enable SSL for an LDAP service, the system does not change the port number from 389 to 636, as is required. To change the port number from the command line, use the port option, for example, ldap [name] ssl enabled port 636.

ssl-ca-cert-file

Specifies the name of an SSL CA certificate using the full path to the file. The default value is none.

ssl-check-peer

Specifies whether the system checks an SSL peer. The default value is disabled.

ssl-ciphers

Specifies SSL ciphers. The default value is none.

ssl-client-cert

Specifies the name of an SSL client certificate. The default value is none.

ssl-client-key

Specifies the name of an SSL client key. The default value is none.

ssl-cname-field

Specifies the value from the client certificate that provides the client name. The client name must appear in either the subjectName or subjectAltName (SAN) fields in the X.509v3 certificate. If it appears in the subjectName field, the client name must be the commonName (CN). If the client name appears in the SAN, it will have the specified type. If san-other is specified, the ssl-cname-otheroid must provide the OID of the UTF8 string containing the client name. The choices are: subjectname-cn, san-other, san-email, san-dns, san-x400, san-dirname, san-ediparty, san-uri, san-ipadd, or san-rid. The default value is subjectname-cn.

ssl-cname-otheroid

Specifies the OID in dotted-decimal format of the UTF8 string in the client's X.509v3 subjectAltName "other" attribute. This value is required when ssl-cname-field is san-other. The default value is none.

sso Enables or disables Single Sign On (SSO) functionality. SSO eliminates the need to administer and maintain multiple user logons and eliminates the need for users to enter their credentials multiple times. When SSO is disabled, the user will be prompted to authenticate into the BIG-IP. The default is off.

user-template

Specifies a user template for the LDAP application to use for authentication. The default value is none.

version

Specifies the version number of the LDAP application. The default value is 3.

warnings

Enables or disables warning messages. The default value is enabled.

SEE ALSO

auth user, create, delete, glob, list, modify, regex, run, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

auth ldap

NAME

ldap - Configures an LDAP configuration object for implementing remote LDAP-based authentication of BIG-IP(r) system users.

MODULE

auth

SYNTAX

Configure the ldap component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

create ldap [name]

modify ldap [name]

options:

bind-dn [[account dn] | none]

bind-pw [none | [password]]

bind-timeout [integer]

check-host-attr [disabled | enabled]

check-roles-group [disabled | enabled]

debug [disabled | enabled]

description [string]

filter [[filter name] | none]

group-dn [[group dn] | none]

group-member-attr [[attribute] | none]

idle-timeout [integer]

ignore-auth-info-unavail [no | yes]

ignore-unknown-user [disabled | enabled]

include [string]

login-attribute [[account name] | none]

port [[name] | [integer]]

referrals [no | yes]

scope [base | one | sub]

search-base-dn [[search base dn] | none]

search-timeout [integer]

servers [add | delete | replace-all-with] {
[[ip address] | [server name] ...] }

servers none

ssl [disabled | enabled]

ssl-ca-cert-file [[file name] | none]

ssl-check-peer [disabled | enabled]

ssl-ciphers [[string] | none]

ssl-client-cert [[string] | none]

ssl-client-key [[string] | none]

user-template [[string] | none]

version [integer]

warnings [disabled | enabled]

edit ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ldap

list ldap [[[name] | [glob] | [regex]] ...]

show running-config ldap

show running-config ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete ldap [name]

DESCRIPTION

LDAP authentication is useful when the BIG-IP system users authentication or authorization data is stored on a remote LDAP server or a Microsoft(r) Windows(r) Active Directory(r) server, and you want the user credentials to be based on basic HTTP authentication (that is, user name and password).

To authenticate BIG-IP system users when their authentication data is stored on a remote LDAP server, you create an LDAP configuration object, and then activate the object.

The following steps describe how to configure LDAP authentication for BIG-IP system users:

1. Use the ldap component in the auth module to configure an LDAP configuration object.

2. To activate LDAP authentication for BIG-IP system users, run the command sequence `modify / auth source type ldap`

EXAMPLES

```
create ldap bigip_ldap_auth servers add {my_ldap_server}
```

Creates a configuration object named `bigip_ldap_auth`

```
delete ldap bigip_ldap_auth
```

Deletes the configuration object named `bigip_ldap_auth`.

OPTIONS

`bind-dn`

Specifies the distinguished name of an account to which to bind to perform searches. This search account is a Read-only account. You can also use the admin account as the search account. If an administrative distinguished name is not specified, then a bind is not attempted. The default value is none.

Note that if the remote server is a Microsoft Windows Active Directory server, the distinguished name must be in the form of an email address.

`bind-pw`

Specifies the password for the search account created on the LDAP server. This option is required if you enter a value for the `bind-dn` option. The default value is none.

`bind-timeout`

Specifies a bind timeout limit, in seconds. The default value is 30.

`check-host-attr`

Confirms the password for the bind distinguished name. This option is optional. The default value is disabled.

`check-roles-group`

Specifies whether to verify a user's group membership given in the remote-role definitions, formatted as `*member*of="group-dn"`. The default value is disabled.

`debug`

Enables or disables syslog-ng debugging information at the LOG DEBUG level. The default value is disabled. F5 Networks does not recommend using this option for normal configuration.

`description`

User defined description.

`filter`

Specifies a filter. Use this option for authorizing client traffic. The default value is none.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`group-dn`

Specifies the group distinguished name. The system uses this option for authorizing client traffic. The default value is none.

`group-member-attribute`

Specifies a group member attribute. The system uses this option for authorizing client traffic. The default value is none.

`idle-timeout`

Specifies the idle timeout, in seconds, for connections. The default value is 3600 seconds.

`ignore-auth-info-unavail`

Specifies whether the system ignores authentication information if it is not available. The default value is no.

`ignore-unknown-user`

Specifies whether the system ignores a user that is unknown. The default value is disabled.

`include`

Specifies configurations to add for `nslcd` conf file.

`login-attribute`

Specifies a logon attribute. Normally, the value for this option is `uid`; however, if the server is a Microsoft Windows Active Directory server, the value must be the account name `samaccountname` (not case-insensitive). The default value is none.

`name` Specifies a unique name for the component. This option is required for the commands `create` and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`port` Specifies the port number or name for the LDAP service. Port 389 is typically used for non-SSL and port 636 is used for an SSL-enabled LDAP service. The default value is `ldap`.

`referrals`

Specifies whether automatic referral chasing should be enabled. The default value is yes.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

scope

Specifies the search scope. The default value is sub. The possible values are:

base The search scope is base object. The base value is almost never useful for name service lookups.

one The search scope is one level.

sub The search scope is a subtree.

search-base-dn

Specifies the search base distinguished name. The default value is none.

search-timeout

Specifies the search timeout, in seconds. The default value is 30.

servers

Specifies the LDAP servers that the system must use to obtain authentication information. You must specify a server when you create an LDAP configuration object.

ssl Enables or disables SSL functionality. The default is disabled.

Note that when you use tmsh to enable SSL for an LDAP service, the system does not change the port number from 389 to 636, as is required. To change the port number from the command line, use the port option, for example, ldap [name] ssl enabled port 636.

ssl-ca-cert-file

Specifies the name of an SSL CA certificate using the full path to the file. The default value is none.

ssl-check-peer

Specifies whether the system checks an SSL peer. The default value is disabled.

ssl-ciphers

Specifies SSL ciphers. The default value is none.

ssl-client-cert

Specifies the name of an SSL client certificate. The default value is none.

ssl-client-key

Specifies the name of an SSL client key. The default value is none.

user-template

Specifies a user template for the LDAP application to use for authentication. The default value is none.

version

Specifies the version number of the LDAP application. The default value is 3.

warnings

Enables or disables warning messages. The default value is enabled.

SEE ALSO

auth user, create, delete, glob, list, modify, regex, run, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2019-12-23 auth ldap(1)

auth login-failures

NAME

login-failures - Displays or resets the status of the accounts of users whose attempts to log in to the BIG-IP(r) system have failed.

MODULE

auth

SYNTAX

Configure the login-failures component within the auth module using the following syntax.

MODIFY

reset-stats login-failures

options:
username

DISPLAY

show login-failures

options:
field-fmt
username

DESCRIPTION

Users assigned a role of Administrator can reset the status of a user who is locked out of the BIG-IP system due to enforcement of a company's security requirements. Users assigned other roles can only view login failures.

EXAMPLES

show login-failures

Displays the login failure status of all users.

show login-failures joe

Displays login failure status for the user joe.

reset-stats login-failures

Resets the failed login counters for all users to zero and unlocks all users.

reset-stats login-failures joe

Resets the failed login counter for the user joe to zero and unlocks the user joe.

OPTIONS

show For information about the options that you can use with the show command, see help show.

username

Specifies a user account to display or reset.

SEE ALSO

auth user, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2012-04-03 auth login-failures(1)

auth partition

NAME

partition - Configures administrative partitions that implement access control for BIG-IP(r) system users.

MODULE

auth

SYNTAX

Configure the partition component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

create partition [name]

modify partition [name]

options:

default-route-domain [ID]

description [string]

DISPLAY

list partition

list partition [[[name] | [glob] | [regex]] ...]

show running-config partition

show running-config partition [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete partition [name]

options:
all

DESCRIPTION

An administrative partition is a logical container that you create, containing a defined set of BIG-IP system objects, such as virtual servers, pools, and profiles. When a specific set of objects resides in a partition, you can then give certain users the authority to view and manage the objects in that partition only, rather than all objects on the BIG-IP system. This gives a finer degree of administrative control.

You can configure administrative partitions, only if the Administrator user role is assigned to your user account.

EXAMPLES

```
create partition partition_A description "Repository for application_A objects"
```

Creates a partition named partition_A that contains objects related to application_A.

```
delete partition partition_B
```

Deletes the partition named partition_B.

OPTIONS

description

Describes the contents of the partition. If you use spaces in the description, you must put quotation marks around the descriptive text, for example, "This partition contains local traffic management objects for managing HTTP traffic."

default-route-domain

Specifies the ID of the route domain that is associated with the IP addresses that reside in the partition. For more information, see help net route-domain.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

auth user, create, delete, glob, list, modify, net route-domain, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-03-21 auth partition(1)

auth password-policy

NAME

password-policy - Specifies the parameters of the valid passwords for the BIG-IP(r) system.

MODULE

auth

SYNTAX

Configure the password-policy component within the auth module using the syntax shown in the following sections.

MODIFY

```
modify password-policy
```

options:

expiration-warning [integer]

max-duration [integer]

max-login-failures [integer]

min-duration [integer]

minimum-length [integer]

password-memory [integer]

policy-enforcement [disabled | enabled]

required-lowercase [integer]

required-numeric [integer]

required-special [integer]

required-uppercase [integer]
lockout-duration [integer]

DISPLAY

list password-policy
list password-policy
show running-config password-policy
show running-config password-policy
options:
 all-properties
 non-default-properties
 one-line

DESCRIPTION

Users assigned a role of Administrator or Resource Administrator can modify a password policy for the BIG-IP system to enforce a company's security requirements by defining the parameters for valid passwords. Users assigned other roles can view password policies.

EXAMPLES

```
password-policy max-duration 90 min-duration 30 minimum-length 6 required-lowercase 2 required-uppercase 2  
required-special 1 required-numeric 1 expiration-warning 5
```

Creates a password policy that specifies that passwords are valid for a maximum of 90 days and a minimum of 30 days. Also specifies that to be valid, a password must contain at least 6 characters, but not more than 10 characters, including 2 lowercase alpha characters, 2 uppercase alpha characters, and 1 number. Additionally, this policy specifies that the system automatically warns users five days before their passwords expire.

```
list password-policy
```

Displays the password policy.

OPTIONS

expiration-warning

Specifies the number of days before a password expires. Based on this value, the BIG-IP system automatically warns users when their password is about to expire. The default value is 7 days.

max-duration

Specifies the maximum number of days a password is valid. The default value is 99999.

max-login-failures

Specifies the number of consecutive unsuccessful login attempts that the system allows before locking out the user. The default value is 0 (zero - disabled).

min-duration

Specifies the minimum number of days a password is valid. The default value is 0 (zero).

minimum-length

Specifies the minimum number of characters in a valid password. The default value is 6.

password-memory

Specifies whether the user has configured the BIG-IP system to remember a password on a specific computer. The default value is 0 (zero).

policy-enforcement

Enables or disables the password policy on the BIG-IP system. The default value is disabled.

required-lowercase

Specifies the number of lowercase alpha characters that must be present in a password for the password to be valid. The default value is 0 (zero).

required-numeric

Specifies the number of numeric characters that must be present in a password for the password to be valid. The default value is 0 (zero).

required-special

Specifies the number of special characters that must be present in a password for the password to be valid. The default value is 0 (zero).

required-uppercase

Specifies the number of uppercase alpha characters that must be present in a password for the password to be valid. The default value is 0 (zero).

lockout-duration

Specifies the amount of time in seconds that a locked-out user must wait before being allowed to log in again unless manually unlocked.

SEE ALSO

auth user, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016. All rights reserved.

auth password

NAME

password - Prompts for modification of a password, and asks for a confirmation of the new password.

MODULE

auth

SYNTAX

Configure the password component within the auth module using the syntax shown in the following sections.

USAGE

modify password

DESCRIPTION

If you are assigned the user role of Administrator or User Manager, you can change another user's password.

For example, from within the auth module, run the following command sequence: modify password [user name].

The system prompts you for a new password for the specified user, and then to confirm the new password.

If you are assigned any other user role, the system prompts you to change your own password, and then confirm the new password.

To change a password from within another module, use the full path to the password.

EXAMPLES

```
(tmos.auth)# modify password
```

From within the auth module, displays the new password: prompt.

```
(tmos.gtm)# modify / auth password
```

From within the gtm module, displays the new password: prompt.

SEE ALSO

auth user, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 auth password(1)

auth radius-server

NAME

radius-server - Configures a RADIUS server for implementing remote RADIUS-based authentication of BIG-IP(r) system users.

MODULE

auth

SYNTAX

Configure the radius-server component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

```
create radius-server [name]
```

```
modify radius-server [name]
```

options:

```
app-service [[string] | none]
```

```
description [string]
```

```
port [ [name] | [number] ]
```

```
secret [none | ["string"] ]
```

```
server [ [hostname] | [IP address] | none]
```

```
timeout [integer]
```

```
edit radius-server [ [ [name] | [glob] | [regex] ] ...]
```

options:

all-properties
non-default-properties

DISPLAY

list radius-server
list radius-server [[[name] | [glob] | [regex]] ...]
show running-config radius-server
show running-config radius-server [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE

delete radius-server [name]

DESCRIPTION

To authenticate BIG-IP system users when their authentication data is stored on a remote RADIUS server, you configure a RADIUS server, configure a RADIUS configuration object that references that RADIUS server, and then activate RADIUS authentication for the BIG-IP system. In this case, client credentials are based on basic HTTP authentication (that is, user name and password).

To configure RADIUS authentication for the BIG-IP system:

1. Use the radius-server component in the auth module to configure a RADIUS server.
2. Use the radius component in the auth module to create a RADIUS configuration object that references the RADIUS server you created in the Step 1. For more information about creating a RADIUS configuration object, see help radius.
3. To activate RADIUS authentication for BIG-IP system users, type the following command sequence: modify / auth source type radius

EXAMPLES

```
create radius-server bigip_auth_radius_server secret "This is the secret." server 10.1.1.1
```

Creates a RADIUS server component named bigip_auth_radius_server.

```
delete radius-server bigip_auth_radius_server
```

Deletes the RADIUS server component named bigip_auth_radius_server.

OPTIONS

app-service

Specifies the name of the application service to which the RADIUS server belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the RADIUS server. Only the application service can modify or delete the RADIUS server.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition in which the radius server resides.

port Specifies the port for RADIUS authentication traffic. The default value is 1812.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Specifies the secret key the system uses to encrypt and decrypt packets sent from or received by the server. This option is required.

server

Specifies the host name or IP address of the RADIUS server. This option is required.

timeout

Specifies the timeout value in seconds. The default value is 3.

SEE ALSO

auth radius, auth user, create, delete, glob, list, modify, regex, run, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

auth radius

NAME

radius - Configures a RADIUS configuration object for implementing remote RADIUS-based authentication of BIG-IP(r) system users.

MODULE

auth

SYNTAX

Configure the radius component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

create radius [name]

modify radius [name]

options:

accounting-bug [disabled | enabled]

app-service [[string] | none]

client-id [none | [string]]

debug [disabled | enabled]

description [string]

retries [integer]

servers [add | delete | replace-all-with]

{ [[hostname] | [ip address] ...] }

servers [default | none]

service-type [default | login | framed | callback-login |

callback-framed | outbound | administrative |

nas-prompt | authenticate-only |

callback-nas-promit | call-check |

callback-administrative]

edit radius [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list radius

list radius [[[name] | [glob] | [regex]] ...]

show running-config radius

show running-config radius [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete radius [name]

DESCRIPTION

To authenticate BIG-IP system users when their authentication data is stored on a remote RADIUS server, you configure a RADIUS server, configure a RADIUS configuration object that references that RADIUS server, and then activate RADIUS authentication for the BIG-IP system. In this case, client credentials are based on basic HTTP authentication (that is, user name and password).

To configure RADIUS authentication for the BIG-IP system:

1. Use the radius-server component in the auth module to configure a RADIUS server. For more information about creating a RADIUS server, see help radius-server.
2. Use the radius component in the auth module to create a RADIUS configuration object that references the RADIUS server you created in Step 1.
3. To activate RADIUS authentication for BIG-IP system users, type the following command sequence: modify / auth source type radius

EXAMPLES

```
create radius bigip_radius_auth servers add {myradiusserver}
```

Creates a RADIUS configuration object named bigip_radius_auth.

```
delete radius bigip_radius_auth
```

Deletes the RADIUS configuration component named bigip_radius_auth.

OPTIONS

accounting-bug

Enables or disables validation of the accounting response vector. This option is necessary only on older servers. The default value is disabled.

app-service

Specifies the name of the application service to which the RADIUS configuration object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the RADIUS configuration object. Only the application service can modify or delete the RADIUS configuration object.

client-id

Sends a NAS-Identifier RADIUS attribute with string bar. If you do not specify a value for this option, the system uses the pluggable authentication module (PAM) service type. You can disable this feature by specifying a blank client ID.

debug

Enables or disables syslog-ng debugging information at the LOG DEBUG level. F5 Networks does not recommend this option for normal use. The default value is disabled.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

retries

Specifies the number of authentication retries that the BIG-IP local traffic management system allows before authentication fails. The default value is 3.

service-type

Specifies the type of service used for the RADIUS server. The default is default, which behaves as authenticate-only.

servers

Specifies the host names or IP addresses of existing RADIUS servers that the BIG-IP system uses to obtain authentication data.

SEE ALSO

auth radius-server, auth user, create, delete, glob, list, modify, regex, run, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 auth radius(1)

auth remote-role

NAME

remote-role - Creates remote role information in a file that an LDAP, Active Directory(r), RADIUS, or TACACS+ server reads to determine the specific access rights to grant to groups of remotely-authenticated users.

MODULE

auth

SYNTAX

Configure the remote-role component within the auth module using the syntax shown in the following sections.

MODIFY

modify remote-role

options:

description [string]

role-info [add | delete | modify | replace-all-with] {

[group-name] {

options:

```

attribute [string]
console [disabled | tmsh]
description [string]
deny [enabled | disabled]
line-order [integer]
role [acceleration-policy-editor | admin | fraud-protection-manager |
application-editor | auditor | certificate-manager |
firewall-manager | guest | irule-manager | manager |
no-access | operator | resource-admin | user-manager |
web-application-security-administrator |
web-application-security-editor | web-application-security-operations-administrator]
user-partition [all | Common | [name] ]
user-partition [%string]
}
}
role-info none

```

DISPLAY

```

list remote-role
show running-config remote-role
options:
  all-properties
  non-default-properties
  one-line

```

DELETE

You cannot delete the remote-role defaults, you can only modify the values of the options.

DESCRIPTION

You can use the remote-role component to grant access to a specific group of remotely-authenticated users without creating a local user account on the BIG-IP(r) system for each user in the group.

Users assigned the role of Administrator can modify remote roles. Users assigned all other roles can view remote roles.

You can use the variable substitution feature to assign access rights for a group of remote users by specifying a text string variable that is preceded by a leading % character for the options attribute, console, role and user-partition. For example, if you define the remote role for the groups DC1 and DC2 as follows:

```

remote-role {
  role info {
    dc1 {
      attribute "F5-LTM-User-Info-1=DC1"
      console %F5-LTM-User-Console
      line-order 1
      role %F5-LTM-User-Role
      user-partition %F5-LTM-User-Partition
    }
    dc2 {
      attribute "F5-LTM-User-Info-1=DC2"
      line-order 2
    }
  }
}

```

The BIG-IP(r) system attempts to match the value of the attribute option, F5-LTM-User-Info-1=DC1, and then pulls the value of the console, role and user-partition options from the other variables.

Note: If a variable includes an incorrect value, the system does not authorize the user. Additionally, if you have not defined the variables, as with the group DC2 above, the system authenticates the user with the following access rights:

```

console = disabled
role = none
user-partition = none

```

EXAMPLES

```

modify remote-role role-info add { my_managers { attribute
"memberOF=cn=BigIPmanagerGroup,cn=users,dc=mydept,dc=mycompany,dc=com" console disabled line-order 1000 role
100 user-partition all } }

```

Configures a remote role, named my_managers, for LDAP authentication, by creating the 1000th line of the /config/bigip/auth/remoterole file, and granting the Manager role (100) in all partitions to the remote users assigned this role.

```

modify remote-role role-info add { my_admins { attribute "NS-Admin-Privilege" console tmsh line-order 1000
role 0 user-partition all } }

```

Configures a remote role, named my_admins, for LDAP authentication, by creating the 2000th line of the /config/bigip/auth/remoterole file, and granting the Administrator role (0) in all partitions to the remote users assigned this role.

```

modify remote-role role-info add { my_managers { attribute "manager_group=manager" console tmsh line-order
3000 user-partition all } }

```

Configures a remote role, named my_managers, for RADIUS or TACACS+ authentication, by creating the 3000th line

of the `/config/bigip/auth/remoterole` file, and granting the Administrator role (0) in all partitions to the remote users assigned this role:

OPTIONS

`description`

Specifies a user-defined description.

`role-info`

Configures the access rights for a specific group of remotely-authenticated users. You can configure the following information for a role:

`attribute`

Specifies an attribute-value pair that an authentication server supplies to the BIG-IP system to match against entries in `/config/bigip/auth/remoterole`. The specified pair typically identifies users with access rights in common. This option is required.

Alternatively, you can use the variable substitution feature (described in the Description section above), and specify a text string variable that is preceded by a leading `%` character.

`console`

Enables or disables console access for the specified group of remotely-authenticated users. The default value is disabled.

When using variable substitution, as described in the Description section of this man page, the variable for the console option must be: `tmsk`.

`deny` Enables or disables remote access for the specified group of remotely-authenticated users. The default value is disabled.

`description`

Specifies a user-defined description.

`group-name`

Specifies the name of the remote role that you are configuring. This option is required.

`line-order`

Specifies the number of the first populated line in the file, `/config/bigip/auth/remoterole`. The LDAP, Active Directory, RADIUS, and TACACS+ servers read this file line by line. The order of the information is important; therefore, F5 Networks recommends that you set the first line at 1000. This allows you, in the future, to insert lines before the first line. This option is required.

`role` Specifies the role that you want to grant to the specified group of remotely-authenticated users. The default value is `no-access`. The available roles are:

`admin`

`fraud-protection-manager`

`application-editor`

`certificate-manager`

`firewall-manager`

`guest`

`manager`

`no-access`

`operator`

`resource-admin`

`web-application-security-administrator`

`web-application-security-editor`

`web-application-security-operations-administrator`

`user-manager`

When using variable substitution, as described in the Description section above, the variable for the role option must evaluate to one of these values: 0 (admin), 20 (resource admin), 40 (user manager), 80 (auditor), 100 (manager), 300 (application editor), 350 (advanced operator), 400 (operator), 450 (firewall manager), 500 (certificate manager), 510 (irule manager), 700 (guest), 800 (web application security administrator), 810 (web application security editor), 820 (web application security operations administrator), 850 (acceleration policy editor), 900 (no-access).

`user-partition`

Specifies the user partition to which you are assigning access to the specified group of remotely-authenticated users. The default value is `Common`. This option is required.

Alternatively, you can use the variable substitution feature (described in the Description section above) and specify a text string variable that is preceded by a leading `%` character.

SEE ALSO

auth remote-user, auth user, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2011, 2013. All rights reserved.

BIG-IP 2019-01-06 auth remote-role(1)

auth remote-user

NAME

remote-user - Configures the default role, partition access, and console access for all remotely authenticated user accounts that have not been added as local user accounts on the BIG-IP(r) system.

MODULE

auth

SYNTAX

Configure the remote-user component within the auth module using the syntax shown in the following sections.

MODIFY

modify remote-user

options:

default-partition [all | Common | [partition name]]
default-role [acceleration-policy-editor | admin |
fraud-protection-manager | application-editor |
auditor | firewall-manager | guest |
irule-manager | manager | no-access |
operator | resource-admin | user-manager |
web-application-security-administrator |
web-application-security-editor | web-application-security-operations-administrator]
description [string]
remote-console-access [disabled | tmsh]

DISPLAY

list remote-user

show running-config remote-user

options:

all-properties
non-default-properties
one-line

DELETE

You cannot delete the remote-user defaults, you can only modify the values of the options.

DESCRIPTION

You can use the remote-user component to configure the default parameters for all the remote user accounts on the BIG-IP system as a group. To assign a different access level to a specific remote user, you must create a local user account for that user on the BIG-IP system. See the auth user man page for more information.

Users assigned the role of Administrator or Resource Administrator can modify the parameters of the remote-user component. Users assigned all other roles can view the parameters of the remote-user component.

EXAMPLES

modify remote-user default-partition Common default-role no access remote-console-access disabled

For all remote users, sets the default partition access to partition Common, the default role to no-access, and the default remote console access to disabled.

modify remote-user default-partition all default-role no access remote-console-access disabled

For all remote users, sets the default partition access to all partitions, the default role to no-access, and the default remote console access to disabled.

OPTIONS

default-partition

Specifies the default partition for all remote user accounts. The default value is all.

default-role

Specifies the default role for all remote user accounts. The default value is no-access.

description

Specifies a user-defined description.

remote-console-access

Specifies whether you are granting this user access to tmsh or disabling remote console access for this

user. The default value is disabled.

SEE ALSO

auth remote-role,auth user, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2018-12-26 auth remote-user(1)

auth source

NAME

source - Configures the authorization source type for a BIG-IP(r) system.

MODULE

auth

SYNTAX

Configure the source component within the auth module using the syntax in the following sections.

MODIFY

modify source

options:

type [active-directory | ldap | local | radius | tacacs | cert-ldap | apm-auth]

fallback [true | false]

DISPLAY

list source

list source [option]

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the source component to set up the authorization source type for the BIG-IP system. The fallback setting enables system auth fallback from remote to local.

EXAMPLES

modify auth source type tacacs

Sets up the authorization source type as tacacs.

list auth source type

Displays the authorization source type.

OPTIONS

type Specifies the default user authorization source. The default value is local. When user accounts that access the system reside on a remote server, the value of this option is the type of server that you are using for authentication, for example, ldap.

fallback

When true and type is set to a remote authentication source, if the remote server is unavailable, authentication will fall back to local authentication. The default value is false.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2018-06-29 auth source(1)

auth tacacs

NAME

tacacs - Configures a TACACS+ configuration object for implementing remote authentication of BIG-IP(r) system users based on TACACS+.

MODULE

auth

SYNTAX

Configure the tacacs component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

create tacacs [name]

modify tacacs [name]

options:

accounting [send-to-first-server | send-to-all-servers]

app-service [[string] | none]

authentication [use-first-server | use-all-servers]

debug [disabled | enabled]

description [string]

encryption [disabled | enabled]

protocol [none | [protocol]]

secret ["[string]"]

servers

[add | delete | replace-all-with] {

[[[hostname[:port]] | [ip address[:port]]] ...]

}

service [[name] | none]

timeout [integer]

edit tacacs [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tacacs

list tacacs [[[name] | [glob] | [regex]] ...]

show running-config tacacs

show running-config tacacs [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete tacacs [name]

DESCRIPTION

To authenticate BIG-IP system users when their authentication data is stored on a remote TACACS+ server, you create a TACACS+ configuration object, and then activate the object.

To configure TACACS+ authentication for BIG-IP system users:

1. Use the tacacs component in the auth module to configure a TACACS+ configuration object.
2. To activate TACACS+ authentication for BIG-IP system users, run the following command sequence: modify / auth source type tacacs

EXAMPLES

```
create tacacs bigip_tacacs_auth servers add {my_tacacs_server}
```

Creates a TACACS+ configuration object named bigip_tacacs_auth.

```
delete tacacs bigip_tacacs_auth
```

Deletes the TACACS+ configuration object named bigip_tacacs_auth.

OPTIONS

accounting

If multiple TACACS+ servers are defined and pluggable authentication module (PAM) session accounting is enabled, sends accounting start and stop packets to the first available server or to all servers. The default value is send-to-first-server.

Possible values are:

send-to-all-servers

The system sends accounting start and stop packets to all servers.

send-to-first-server

The system sends accounting start and stop packets to the first available server.

app-service

Specifies the name of the application service to which the TACACS+ configuration object belongs. The

default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the TACACS+ configuration object. Only the application service can modify or delete the TACACS+ configuration object.

authentication

Specifies the process the system employs when sending authentication requests. The default value is use-first-server.

Possible values are:

use-all-servers

The system sends an authentication request to each server until authentication succeeds, or until the system has sent a request to all servers in the list.

use-first-server

The system sends authentication requests to only the first server in the list.

debug

Enables syslog-ng debugging information at the LOG DEBUG level. F5 Networks does not recommend this option for normal use. The default value is disabled.

description

User defined description.

encryption

Enables or disables encryption of TACACS+ packets. F5 Networks recommends this option for normal use. The default value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

protocol

Specifies the protocol associated with the value specified in the service option, which is a subset of the associated service being used for client authorization or system accounting.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Sets the secret key used to encrypt and decrypt packets sent or received from the server. This option is required.

servers

Specifies the host name or IPv4 address of the TACACS+ server. For each server, a port may optionally be specified in the format hostname:port or IPv4:port. If no port is specified, the default port 49 is used. This option is required.

service

Specifies the name of the service that the user is requesting to be authenticated to use. Identifying the service enables the TACACS+ server to behave differently for different types of authentication requests. This option is required.

timeout Specifies the connect timeout in seconds. The default value is 10 seconds. Zero indicates no timeout.

SEE ALSO

auth user, create, delete, edit, glob, list, modify, regex, run, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2016-08-10 auth tacacs(1)

auth user

NAME

user - Configures user accounts for the BIG-IP(r) system.

MODULE

auth

SYNTAX

Modify the user component within the auth module using the syntax shown in the following sections.

CREATE/MODIFY

```
create user [name]
modify user [name]
options:
  description [text...]
  partition-access [add | modify | delete |replace-all-with { [partition-name] { role [role-name] } } ]
  password [text]
  prompt-for-password
  shell [name]
```

where [role-name]: [acceleration-policy-editor | admin | fraud-protection-manager | application-editor | auditor | certificate-manager | firewall-manager | guest | irule-manager | manager | no-access | operator | resource-admin | user-manager | web-application-security-administrator | web-application-security-editor | web-application-security-operations-administrator]

DISPLAY

```
list user
list user [ [ [name] | [glob] | [regex] ] ... ]
show running-config user
show running-config user [ [ [name] | [glob] | [regex] ] ... ]
options:
  encrypted-password
  one-line
  partition
show user
options:
  field-fmt
```

DELETE

```
delete user [name]
```

DESCRIPTION

You can create user accounts where the user names differ only by case-sensitivity (for example, david and DAVID).

You can configure the partition-access property to grant a user access to more than one partition on the system. In the case where you do not grant the user access to all partitions, you can assign the user a different user role for each partition. A user can have only one role per partition. Any user with a role of Administrator, Resource Administrator, Web Application Security Administrator, or Auditor always has access to all partitions and can have no other role on the system.

Only users with the Administrator or User Manager roles are allowed to create or modify user accounts.

Additionally, only users with the Administrator, Resource Administrator, or User Manager user role can view all of the user accounts in all of the partitions to which the user has access. Therefore, if you have a user role other than one of these roles, you can only view your own user account.

EXAMPLES

```
create user nwinters partition-access add { all-partitions { role guest } }
```

Creates a new user named nwinters with a role of Guest in all partitions.

```
create user tknox password aBcD007 partition-access add { partition1 { role operator } }
```

Creates a new user named tknox with a role of operator in partition named partition1 and sets the user's login password.

```
list user
```

Displays the viewable properties of all user accounts.

```
show user
```

Displays each user role and the corresponding partition access that is currently assigned to the user.

OPTIONS

description
Describes the user account in free form text.

encrypted-password
Displays the encrypted password for the user account.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

Note: User account names are case-sensitive.

partition

Displays the name of the administrative partition in which the user account resides.

partition-access

Specifies the administrative partitions to which the user currently has access. Note that in addition to these partitions, the user also has read access to the shared partitions Common and Root. An exception to this is any user with the role No Access.

role Specifies the user role that pertains to the partition specified by the partition-access property. If you do not want to assign a user role to the user account, specify the value no-access. This prevents the user from accessing the system.

password

Sets the user password during creation or modification of a user account without prompting or confirmation. May not be used with prompt-for-password. Passwords are hidden in log and history files.

prompt-for-password

Indicates that when the account is created or modified, the BIG-IP system prompts the administrator or user manager for both a password and a password confirmation for the account.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

shell

Specifies the shell to which the user has access. Valid values are:

bash Provides an unrestricted system prompt. You can assign access to the bash shell only to users with the Administrator role.

none Specifies no shell access. The user must use the Configuration utility.

tmsh Provides access to the Traffic Management shell. Resource Administrator user role can use the tcpdump, ssldump, or qkview utilities within tmsh shell (run /util). Other user roles may be given this shell, as appropriate.

SEE ALSO

auth partition, auth password, create, delete, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2016. All rights reserved.

BIG-IP 2018-12-26 auth user(1)

cli

cli admin-partitions

NAME

admin-partitions - Set the administrative partition for a BIG-IP(r) configuration file.

MODULE

cli

SYNTAX

Configure the admin-partitions component within the cli module using the syntax in the following sections.

MODIFY

modify admin-partitions

options:

update-partition [name]

DESCRIPTION

You can use the admin-partitions component to set the administrative partition in which configuration will be loaded when a configuration file is being loaded.

This component is only available from a configuration file that is being loaded via the sys config component with the file option.

EXAMPLES

```
cli admin-partitions { update-partition partition_A }
```

Sets the administrative partition in which configuration will be loaded. Configuration that follows this directive will be place in partition_A.

OPTIONS

update-partition

Sets the administrative partition in which you can configure objects.

SEE ALSO

load, sys config, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013. All rights reserved.

BIG-IP 2017-06-26 cli admin-partitions(1)

cli alias private

NAME

private - Configures a user private alias.

MODULE

cli alias

SYNTAX

Configure the alias component within the cli alias module using the syntax in the following sections.

CREATE/MODIFY

create private [name]

options:

command [commandSyntax]

command ["command syntax"]

command "[command syntax]; [command syntax]; ..."

app-service [[string] | none]

description [string]

edit private [name]

options:

all

modify alias [name]

options:

command [commandSyntax]

command ["command syntax"]

command "[command syntax]; [command syntax]; ..."

DISPLAY

list private

list alias [[[name] | [glob] | [regex]] ...]

show running-config private

show running-config private [[[name] | [glob] | [regex]] ...]

options:

all-properties

one-line

non-default-properties

DELETE

delete private [all | [name ... name]]

DESCRIPTION

You can use the private component to create a shortcut that runs a tmsh command sequence. The name of the private alias is what you type on the command line to run the command. If the command sequence for which you are creating an alias contains spaces, it must be enclosed in quotation marks. Command aliases are not case-sensitive.

You can create a private alias that runs multiple commands by entering the command sequences separated by semi-colons.

Private aliases can be used only by the user who created them.

When a batch mode transaction is active, commands that operate on the private component are run immediately and are not added to the transaction.

For more information about aliases, see the Traffic Management Shell (tmsh) Reference Guide.

EXAMPLES

```
create private save command "save config"
```

Creates an alias that saves the running configuration in the stored configuration files from anywhere within

tmsh.

create private stats command "show /sys traffic"

Creates an alias that displays traffic statistics from anywhere within tmsh.

create private nodemonitor command "list /ltm node; list /ltm monitor"

Creates an alias that displays the Local Traffic Manager nodes and monitors.

create private myalias command "show /sys provision ; show /sys license"

Creates an alias that displays license and provisioning information.

create private ltm pool command "list /ltm pool"

Creates an alias that displays the Local Traffic Manager pools from anywhere within tmsh.

OPTIONS

command syntax

Specifies the command for which you are creating an alias. To create an alias that runs multiple commands, enter the command sequences separated by semi-colons.

app-service

Specifies the name of the application service to which the alias belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the alias. Only the application service can modify or delete the alias.

description

Specifies the purpose of the alias. If you enable cli preference show-aliases, tmsh displays the description in context-sensitive help (?).

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a name for the alias. This is what you type in tmsh to run the command for which you are creating an alias.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, shared, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-12-20 cli alias private(1)

cli alias shared

NAME

shared - Configures a shared alias.

MODULE

cli alias

SYNTAX

Configure the shared alias component within the cli alias module using the syntax in the following sections.

CREATE/MODIFY

create shared [name]

options:

command [commandSyntax]

command ["command syntax"]

command "[command syntax]; [command syntax]; ..."

app-service [[string] | none]

description [string]

edit shared [name]

options:

all

modify alias [name]
options:
 command [commandSyntax]
 command ["command syntax"]
 command "[command syntax]; [command syntax]; ..."

DISPLAY
list shared
list alias [[name] | [glob] | [regex]] ...]
show running-config shared
show running-config shared [[name] | [glob] | [regex]] ...]
options:
 all-properties
 one-line
 non-default-properties

DELETE
delete shared [all | [name ... name]]

DESCRIPTION

You can use the shared component to create a shortcut to run a tmsh command sequence. The name of the shared alias is what you type on the command line to run the command. If the command sequence for which you are creating an alias contains spaces, it must be enclosed in quotation marks. Command aliases are not case-sensitive.

You can create a shared alias that runs multiple commands by entering the command sequences separated by semi-colons.

Shared aliases can be used by all users.

When a batch mode transaction is active, commands that operate on the shared component are run immediately and are not added to the transaction.

For more information about aliases, see the Traffic Management Shell (tmsh) Reference Guide.

EXAMPLES

create shared save command "save config"

Creates an alias that saves the running configuration in the stored configuration files from anywhere within tmsh.

create shared stats command "show /sys traffic"

Creates an alias that displays traffic statistics from anywhere within tmsh.

create shared nodemonitor command "list /ltm node; list /ltm monitor"

Creates an alias that displays the Local Traffic Manager nodes and monitors.

create shared myalias command "show /sys provision ; show /sys license"

Creates an alias that displays license and provisioning information.

create shared ltmpool command "list /ltm pool"

Creates an alias that displays the Local Traffic Manager pools from anywhere within tmsh.

OPTIONS

command syntax
Specifies the command for which you are creating an alias. To create an alias that runs multiple commands, enter the command sequences separated by semi-colons.

app-service
Specifies the name of the application service to which the alias belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the alias. Only the application service can modify or delete the alias.

description
Specifies the purpose of the alias. If you enable cli preference show-alias, tmsh displays the description in context-sensitive help (?).

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a name for the alias. This is what you type in tmsh to run the command for which you are creating an alias.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, shared, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2015-12-28 cli alias shared(1)

cli global-settings

NAME

global-settings - Configures settings for tmsh

MODULE

cli

SYNTAX

Configure the global-settings component within the cli module using the syntax shown in the following sections.

MODIFY

edit global-settings

options:

all-properties

non-default-properties

modify global-settings

options:

audit [disabled | enabled]

description [string]

idle-timeout [disabled | integer]

scf-backup-number [integer]

service [number | name]

DISPLAY

list global-settings

list global-settings [option]

options:

all-properties

non-default-properties

one-line

DELETE

You cannot delete the default global settings.

DESCRIPTION

You can use the global-settings component to configure multiple settings for tmsh.

EXAMPLES

modify global-settings audit enabled

Enables auditing for tmsh.

modify global-settings idle-timeout 15

Sets the user idle timeout from tmsh to 15 minutes.

OPTIONS

audit

Specifies the global audit level for tmsh. The audited commands are stored in `/var/log/audit`. The default value is enabled. The audit levels are:

disabled

tmsh does not log commands that users enter.

enabled

tmsh audits only commands that users enter. Note that the system does not audit the commands that the command load runs.

description

User defined description.

idle-timeout

If not disabled, log a user in tmsh interactive mode out automatically after a specified set of minutes.

An administrator may change the timeout value at any time and the new policy will take place immediately.

scf-backup-number

Specifies the number of backup single configuration files that the system stores when you enter the following command sequence in tmsh:

load sys config file

When you run the command, the system saves the single configuration file. By default, the system saves two backup single configuration files. For example, if you set the scf-backup-number option to 3, after you run the command sequence tmsh load sys config file for the third time, the system has three versions of the single configuration file: /var/local/scf/backup.scf, /var/local/scf/backup-1.scf, and /var/local/scf/backup-2.scf. The newest file is /var/local/scf/backup.scf.

service

Specifies the format in which tmsh displays a service. The default value is name. The options are:

name Displays a service using a protocol name, for example, http.

number

Displays a service using a numeric value, for example, 192.168.10.20:80, where 80 indicates http.

SEE ALSO

edit, list, modify, run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016 All rights reserved.

BIG-IP 2016-03-14 cli global-settings(1)

cli history

NAME

history - Displays a list of commands in the order in which you ran the commands.

MODULE

cli

SYNTAX

Use the history component within the cli module to display a numbered list of commands in the order the commands were issued.

DISPLAY

```
show history
!  
!!  
![string]
```

DESCRIPTION

You can use the history component to display a numbered list of the commands that you have run in tmsh. The commands display in the order in which you ran the commands, and each command is identified by an entry ID. The larger the entry ID of the command, the more recently you ran the command.

To rerun a command from the history list, type q to close the list and return to the tmsh prompt, and then enter an exclamation point (!) followed by the entry ID of the command that you want to run.

EXAMPLES

```
!  
  
show history  
  
Either of the two previous commands, displays the command history list.
```

```
!5  
  
Runs the fifth command in the command history list.
```

```
!!  
  
Runs the previously issued command.
```

```
!create  
  
Runs the last command that begins with create.
```

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-11 cli history(1)

cli preference

NAME

preference - Configures tmsh preferences.

MODULE

cli

SYNTAX

Configure the preference component within the cli module using the syntax shown in the following sections.

MODIFY

edit preference

modify preference [option]

options:

alias-path [string list]
app-service [[string] | none]
confirm-edit [disabled | enabled]
display-threshold [integer]
editor [nano | vi]
history-date-time [disabled | enabled]
history-file-size [integer]
history-size [integer]
keymap [default | emacs | vi]
list-all-properties [disabled | enabled]
pager [disabled | enabled]
prompt { [avc-count config-sync-status current-folder
fully-qualified-host host mcp-load-status
mcp-state multi-line status user user-role] | none }
show-aliases [disabled | enabled]
stat-units [default | exa | gig | kil | meg | peta | raw |
tera | yotta | zetta]
suppress-warnings [all | config-version | none]
table-indent-width [integer]
tcl-syntax-highlighting [disabled | enabled]
video [disabled | enabled]
warn [bell | disabled | visual-bell]

edit preference

options:

all-properties

DISPLAY

list preference

list preference [option]

show running-config preference

show running-config preference [option]

options:

all-properties
one-line

DESCRIPTION

You can use the preference component to configure tmsh to meet your specific needs.

EXAMPLES

```
modify preference display-threshold 500
```

Configures tmsh to retrieve up to 500 objects before requiring a user response to the question, "Display all items? (y/n)."

```
modify preference history-file-size 80
```

Configures the maximum number of commands that a user can view in the command history list to be 80.

```
modify preference history-size 1000
```

Configures the maximum number of commands that tmsh saves in a user's .tmsh_history file to be 1000 commands.

```
modify preference suppress-warnings config-version
```

Configures tmsh to suppress warning messages for configuration version related (for backward compatibility of configuration).

OPTIONS

alias-path

Specifies the search paths for shared aliases. The shared aliases could be in multiple locations, only ones on the search paths can be used. If a folder is deleted from the system it will be automatically remove from the alias-path.

app-service

Specifies the name of the application service to which the preference belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the preference. Only the application service can modify or delete the preference.

confirm-edit

Specifies whether the command edit prompts for confirmation before saving changes. The default value is enabled.

Note that the value of this option does not affect the behavior of the editor if the changes made in the editor result in a failed update. In this case, tmsh always prompts the user to either re-edit the file or discard the changes. The options are:

enabled

tmsh prompts a user to either submit (y), discard (n), or edit (e) the changes made to a component within the editor.

disabled

tmsh does not prompt the user, but instead, immediately submits the changes made in the editor.

display-threshold

Specifies the maximum number of objects that tmsh displays without requiring a user response to the question, "Display all [number] items? (y/n)." You can specify from 0 (zero) through 4,294,967,265 objects. If you set this option to 0 (zero), tmsh displays an unlimited number of objects without requesting a response.

editor

Specifies the editor that the command edit invokes. Users assigned the user role of Administrator can select nano or vi. Users assigned other user roles must use nano.

history-date-time

Specifies whether tmsh displays in the command history the date and time that each command was issued. The default value is disabled.

Note that the command history file, `~/.tmsh-history-[user]`, always contains the date and time that a command was issued.

history-file-size

Specifies the maximum number of tmsh commands that the system saves in each user's `.tmsh_history` file. If you set this option to 0 (zero), the system does not save tmsh commands in the file. The maximum value is 100,000. For performance reasons, the system does not truncate the file after a user enters a command. Instead, the system truncates the file after a user exits tmsh.

history-size

Specifies the number of commands that a user can view or search in the command history list. The maximum number of commands is 100,000. The default value is 500.

If you set this option to 0 (zero), the system does not add commands to the list of commands in memory; however, the system does write commands to the `.tmsh_history` file, unless the `history-file-size` option is set to 0 (zero).

When you change the value of this option, the system renumbers the commands listed in memory; however, the commands remain in the same order.

keymap

Specifies the keyboard bindings that you want tmsh to use. The default value is default. The options are default, emacs, and vi.

list-all-properties

Specifies whether the system displays all of the properties of a component by default when you run the command list. The default value is disabled.

pager

Specifies whether the system sends the output of the tmsh commands list and show to less. The default value is enabled.

prompt

Specifies the information that you want to display in the tmsh prompt. By default the prompt displays `user_name@host_name(tmos-current_module)#`. The options are:

avc-count

Displays the current SELinux Access Vector Cache in the tmsh prompt. The value displayed in the prompt indicates the number of times SELinux has denied access to a protected resource. The default is to not display this information.

config-sync-status

Displays global sync status in the tmsh prompt. The status displayed in the prompt indicates the rolled-up sync status of all the device groups where the local device resides. The default is to display this information.

current-folder

Displays the current working folder in the tmsh prompt. The default is to not display this information.

fully-qualified-host
Displays the fully qualified host name in the tmsh prompt. The default is to not display this information.

host Displays the host name in the tmsh prompt. The default is to display the host name in the prompt.

mcp-load-status
Displays the configuration file load status in the tmsh prompt. This information is also available in the Last Configuration Load Status of the show sys mcp command output. The default is to not display this information.

mcp-state
Displays the running phase of the mcpd service in the tmsh prompt. This information is also available in the Running Phase of the show sys mcp command output. The default is to not display this information.

multi-line
Displays the tmsh prompt on multiple lines, with information on the first line, and a pound sign (#) on the second line, for example:

```
(Common:all) operator1@6400(tmos.cli)
```

```
#
```

The multi-line option is disabled by default.

none Sets the tmsh prompt to display (tmos.current_module)#, where the system replaces current_module with the name of the module within which you are working.

status
Displays the system status in the tmsh prompt. The default is to display system status in the prompt.

user Displays the user name in the tmsh prompt. The default value is to display the user name in the prompt.

user-role
Display the user's current role in the tmsh prompt. By default the user role is not displayed in the prompt.

show-aliases
Specifies whether the system displays aliases in the results of the command completion and context-sensitive help features. The default value is enabled.

suppress-warnings
Specifies the type of warning messages which needs to be suppressed. The default value is none.

stat-units
Specifies the default unit in which the system displays statistics. The options are:

default
Displays data in the simplest units. For example, if the value of the data is 1,200,001, the system displays 1.20M; however, if the value of the data is 1,200, the system displays 1.2K.

exa Display data in parts per quintillion.

gig Displays data in parts per billion.

kil Displays data in parts per thousand.

meg Displays data in parts per million.

peta Displays data in parts per quadrillion.

raw Displays raw data.

tera Displays data in parts per trillion.

yotta
Displays data in parts per septillion.

zetta
Displays data in parts per sextillion.

table-indent-width
Specifies the indent width when tmsh displays the child object tables in a show command. You can specify from 0 (zero) through 10. If you set this option to 0 (zero), tmsh displays child object tables without any indent.

tcl-syntax-highlighting
Specifies whether Tcl syntax highlighting will be enabled in the editor. This setting only applies if your editor preference is set to vi. The default value is disabled.

video

Enables or disables any video features used to highlight text. The default value is enabled.

warn Specifies how the system warns you when you make an incorrect keystroke. The default value is bell.

The options are:

bell Sounds a bell.

disabled
Disables the warning function.

visual-bell
Displays a visual warning.

SEE ALSO

edit, list, modify, show, sys mcp-state, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 cli preference(1)

cli script

NAME

script - Automates tmsh using Tool Command Language (Tcl).

MODULE

cli

SYNTAX

Configure the script component within the cli module using the syntax shown in the following sections.

EDIT

```
create script [name]
modify script [name]
options:
  app-service [[string] | none]
  description [string]
  ignore-verification [true | false]
  script-checksum [[string] | none]
  script-signature [[string] | none]
```

```
edit script [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
```

DISPLAY

```
list script
list script [ [ [name] | [glob] | [regex] ] ... ]
show running-config script
show running-config script [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
```

DELETE

```
delete script [name]
```

GENERATE

Note: generate cryptographic signature or checksum based on cli script text.

```
generate cli script [name]
```

```
options:
  checksum
  signature
```

RUN

```
run script [name] [options ...]
options:
  file [file name] [options ...]
  verbatim-arguments [file option] [file name] [options ...]
```

The options that are available depend on which script you are running.

The file option is limited to users with the role of administrator.

DESCRIPTION

You can use the script component to build Tcl scripts to automate management of the BIG-IP(r) system. By combining command aliases with scripts, you can extend tmsh to build commands that are customized to your environment.

To do this, place the content of the script inside one or more Tcl procedures. The content of a script cannot exceed 65,000 bytes. However, a script can include other scripts. For more information about including scripts in other scripts, see `tmsh::include` following.

Starting with BIGIP 11.5.0, tmsh commands are versioned. The tmsh active version should be specified in scripts. This will avoid breaking scripts due to changes in tmsh syntax in the different versions. See examples below for how to use it in a script. Without tmsh active version specified, scripts will run on the current active version. By default, the active version will be the latest cli version.

You can use the following procedures in the manner specified:

`script::run`

tmsh invokes the procedure `script::run` when you issue the command sequence `run / cli script [name]`. A script is run relative to the module in which the run command is invoked.

The `script::run` procedure must be defined in the script named by the run command. Scripts that are included by `tmsh::include` are not required to implement the procedure `script::run`.

`script::help`

Provides context sensitive help. A script is not required to implement `script::help`.

`script::tabc`

Provides context sensitive help. A script is not required to implement `script::tabc`.

`script::init`

tmsh calls the procedure `script::init` before calling one of the following procedures: `script::run`, `script::help`, or `script::tabc`. The `script::init` procedure can use the Tcl variable `tmsh::csh` to determine which one of these three procedures tmsh invokes after `tmsh::init`.

Additionally, you can use the procedure `script::init` to initialize global variables. A script is not required to implement `script::init`.

EXAMPLES

```
edit script myscript
```

Creates or modifies the script myscript.

```
edit script myscript yourscript
```

Creates or modifies the scripts myscript and yourscript at the same time.

```
list script myscript
```

Displays the contents of the script myscript.

```
delete script [name]
```

Deletes the script myscript from the system.

```
run script myscript [arguments ...]
```

Runs the script myscript. The system passes arguments to the script in the following Tcl variables:

• `tmsh::argc` contains the number of arguments including the name of the script.

• `tmsh::argv` contains the list of argument values. The first item in `tmsh::argv` is always the name of the script.

Tip: You can create an alias for the command sequence `run / cli script [name]` using the cli alias component. For more information, see `help cli alias`.

```
run script verbatim-arguments myscript [arguments ...]
```

Runs the same commands as `run script myscript [arguments...]` above, except the system passes all arguments specified in the command as one argument to the script. Note that you do not need to enclose the argument list in double quotes, and you do not need to escape special characters.

```
generate my_script checksum
```

Generate a checksum for the script text and add the checksum as a property.

```
generate my_script signature signing-key my_key
```

Generate a signature for the script text using the specified private key and add the signature as a property.

Note: For a script which includes a checksum or signature to successfully load, the script text contents must match the stored checksum or signature. To temporarily stop the verification of signature or checksum and still retain the checksum or signature, the ignore-verification attribute must be set to true. This is done by editing the script and adding the ignore-verification attribute.

To completely clear the signature or checksum, simply set the attribute `script-signature` or `script-checksum` to

empty string "". By doing so, the script will be processed as if it was never signed or checksummed.

```
modify script /Common/my_script { proc script::init {} { }
```

```
proc script::run {} { }
```

```
proc script::help {} { }
```

```
proc script::tabc {} { }  
ignore-verification true  
script-checksum 74778e7b13016e0b9329a17f8d2da601  
total-signing-status checksum  
verification-status checksum-verified }
```

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum

Generate a checksum for the script text and add the checksum to the script as a property. Only for use with the generate command.

description

A user defined description.

file Specifies that the script to be run should come from a file located on the file system rather than a script from the configuration.

glob Displays the scripts that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the script. This option is required for the edit and delete commands.

regex

Displays the scripts that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

signature

Generate a signature for the script text using the specified private key and add the signature to the script as a property. Only for use with the generate command.

signing-key

The private key to use for signing the script. Only for use with the signature option.

verbatim-arguments

Specifies that the arguments at the end of the command should not be tokenized by tmsh prior to being sent to the script. This is useful when the script is wrapping another utility that takes arguments.

CONFIGURATION AND STATUS ACCESSORS

The following Tcl commands mirror tmsh commands. For example, the Tcl tmsh::create command accepts the same components, object identifiers, and properties that the tmsh create command accepts.

tmsh::cd [args...]

Runs the cd command using the specified arguments.

tmsh::cp [args...]

Runs the cp command using the specified arguments.

tmsh::create [args...]

Runs the create command using the specified arguments.

tmsh::delete [args...]

Runs the delete command using the specified arguments.

tmsh::install [args...]

Runs the install command using the specified arguments.

tmsh::generate [args...]

Runs the generate command using the specified arguments.

tmsh::list [args...]

Runs the list command using the specified arguments. The system returns the results as a string.

tmsh::load [args...]

Runs the load command using the specified arguments.

tmsh::modify [args...]

Runs the modify command using the specified arguments.

tmsh::publish [args...]

Runs the publish command using the specified arguments.

tmsh::pwd

Runs the pwd command.

`tmsh::reset-stats [args...]`

Runs the reset-stats command using the specified arguments.

`tmsh::restart [args...]`

Runs the command restart using the specified arguments.

`tmsh::run [args...]`

Runs the run command using the specified arguments.

`tmsh::save [args...]`

Runs the save command using the specified arguments.

`tmsh::show [args...]`

Runs the show command using the specified arguments. The system returns the results as a string.

`tmsh::start [args...]`

Runs the command start using the specified arguments.

`tmsh::stop [args...]`

Runs the command stop using the specified arguments.

The following Tcl commands provide structured access for retrieving configuration, statistics, and status information.

`tmsh::get_config [args...]`

Returns a list of Tcl objects. Each of these objects can be passed to the commands that accept an `$obj` argument. The arguments for this command are the same as for the `tmsh list` command.

`tmsh::get_status [component] [args...]`

Returns a list of Tcl objects that can be passed to the following commands that accept an `$obj` argument. The arguments to this command are the same as the `tmsh show` command.

This command can only be used on components that accept the `field-fmt` option. The `field-fmt` option is automatically appended to the argument list. The `tmsh help` pages identify if a component supports the `field-fmt` option.

That there are very few components that have status and statistics that do not support the `field-fmt` option, and in those cases you can use the Tcl `tmsh::show` command to retrieve the object in the form of a Tcl string object.

A component must be specified, for example, `tmsh::get_status ltm pool`.

`tmsh::get_type $obj`

Returns the type identifier associated with the object. The `$obj` argument must be an object that was returned by either of the Tcl `tmsh::get_config` or `tmsh::get_status` commands.

`tmsh::get_name $obj`

Returns the object identifier associated with the object. The `$obj` argument must be an object that was returned by either of the Tcl commands `tmsh::get_config` or `tmsh::get_status`.

`tmsh::get_field_names [value | nested] $obj`

Returns a list of field names (not the value associated with a field) that are present in an object. The value fields are simple values or lists (for example, an integer or a string). The nested fields are a collection of zero or more nested objects, where the nested objects have their own fields (for example, pool members, and virtual server profiles).

The `$obj` argument must be an object that was returned by the Tcl `tmsh::get_config` or `tmsh::get_status` commands. If the object was retrieved using the Tcl `tmsh::get_config` command, the field names are identical to those that are displayed by the `tmsh list` command. If the object was retrieved using the Tcl `tmsh::get_status` command, the fields are identical to those that the system displays using the `tmsh show` command with the `field-fmt` option.

`tmsh::get_field_value $obj [field name] [Tcl variable]`

Retrieves the value of field name.

The Tcl variable is optional. The behavior of this command depends on whether field name is present in `$obj` and a Tcl variable is present in the command.

• If field name is present in `$obj`, and a Tcl variable is present, the Tcl variable is set to the value of field name and the command returns 1.

• If field name is not present in `$obj`, and a Tcl variable is present, the command returns 0 (zero).

• If field name is present in `$obj`, and a Tcl variable is not present, the command returns the field value.

• If field name is not present in `$obj`, and a Tcl variable is not present, the command raises an error that causes the script to stop. You can use the Tcl command `catch` to recognize the error and continue to run the script.

The `$obj` argument must be an object that was returned by the Tcl `tmsh::get_config` or `tmsh::get_status` commands, or a nested object obtained from the Tcl `tmsh::get_field_value` command.

If the field is a set of nested objects, the Tcl object that the system returns is a list of objects, where each of the objects can contain fields. Each of the objects can be passed to the Tcl `tmsh::get_field_value` command. If the field is not a nested object the system returns a single Tcl string

object.

TRANSACTION CONTROL

The following Tcl commands are specific to the tmsh Tcl API. There are no corresponding commands available in tmsh.

`tmsh::begin_transaction`

Begins an update transaction. The Tcl `tmsh::create`, `tmsh::delete`, and `tmsh::modify` commands that are issued before the next Tcl `tmsh::commit_transaction` command are submitted as a single update.

The system rolls back all of the commands if any of the commands fail.

`tmsh::commit_transaction`

Runs the commands that have been issued since the last Tcl `tmsh::begin_transaction` command. The system validates all of the commands against the running configuration. If any one of the commands fail, the system does not apply any of the commands to the running configuration.

`tmsh::cancel_transaction`

Cancels all commands that you have issued since the last Tcl `tmsh::begin_transaction` command.

Important: You cannot use these Tcl commands inside an active transaction:

Â· `tmsh::list`

Â· `tmsh::show`

Â· `tmsh::get_config`

Â· `tmsh::get_status`

LOGGING

You can use the following Tcl commands to generate log events. These commands affect the behavior of the script and do not affect tmsh. These commands are available only to users who have been assigned either the Administrator or Resource Administrator role.

`tmsh::log_dest [screen | file]`

Specifies whether the system sends events to the screen or to log files. If file is selected, log messages will be directed to `/var/log/ltn`.

`tmsh::log_level [level]`

Specifies the default severity level. The system does not log events below the specified level. The options, listed in decreasing order of severity, are:

Â· emerg

Â· alert

Â· crit

Â· err

Â· warning

Â· notice

Â· info

Â· debug

`tmsh::log [level] "message..."`

Logs the specified message. The level parameter is optional. The level can be one of those described in the Tcl `tmsh::log_level` command.

CUSTOM ISTATS

Custom counter, gauge, and string fields may be created, modified, and retrieved using iRules or tmsh scripts. These custom fields are created on first write and do not need to be declared separately.

Each custom field has a "key" that can be associated with a tmsh configuration object. This key is composed of a tmsh component dotted path, a specific object name or ID, the field type, and the field name. The entire key must be enclosed in quotes.

For example, `"ltn.pool /Common/my_pool counter num_hits"` refers to the `num_hits` counter associated with the LTM pool named `my_pool`, located in the `Common` folder.

These custom fields are displayed with the tmsh `show` command on the associated object.

`istats::incr [key] [amount]`

Increments a custom counter by amount.

`istats::set [key] [value]`

Sets a custom gauge or string to value. Setting a counter to an exact value will only set it in the local segment, but `istats::get` will always read the aggregated (not local) value.

`istats::get [key]`

Returns the latest aggregated value of the custom field or 0 (zero) if it does not exist ("" for string fields).

istats::remove [key]

Removes the custom field from all segments on all blades. Effectively resets a counter to 0.

UTILITIES

The following commands are TCL utility commands.

tmsh::clear_screen

Clears the screen and places the cursor at the upper left of the screen.

tmsh::display [variable | command output]

Provides access to the tmsh pager. Output generated with the Tcl puts command is not paged.

tmsh::display_threshold [integer]

When a script is run, the system disables the option cli preference display-threshold.

You can use the Tcl tmsh::display_threshold command to re-enable the threshold. Re-enabling the threshold in this way causes the script to generate a prompt if you issue the tmsh::list, tmsh::show, tmsh::get_config, or tmsh::get_status commands, and the output that is generated exceeds the threshold. See help cli preference for a description of this option and valid ranges for its value.

tmsh::expand_macro [macro_text] options...

Expands a macro and returns the resulting string. A macro is a string containing macro syntax which can be used for parameter substitution, script and iRule templating, etc. The Macro Syntax includes the following delimiters:

<% The beginning of an expansion code block.

<%= The beginning of an expansion code block. Spool the output after evaluating.

<%D[0-9][0-9] The beginning of a debug/logging code block with the debug threshold set to 0 thru 99.

<%D[0-9][0-9]= The beginning of a debug/logging code block with the debug threshold set to 0 thru 99. Spool the output after evaluating.

%> The end of the current block (works for all types).

Typically, the result of the expand_macro command is used as the input to another command (eg. Irm rule create). The command can be called multiple times within an iApp implementation to expand multiple macros.

macro_text is the blob of text to expand. If not specified, the command will expand the Macro section of the iApp. If no macro_text argument is specified and no Macro section exists for the iApp, an error will be issued.

-vars name_value_pair_list

Specifies a list of additional variables (name/value pairs) which can be referenced within the macro and expanded by the command. All APL variables are automatically available from within the macro, so the -vars option allows a way to specify additional variables from the iApp Implementation section. Since the variables are defined within a Tcl list the format is: { name1 value1 name2 ... nameN valueN }

-debug debug_levels

Specifies a single debug level or list of debug levels for controlling which debug messages get rendered in the expanded output.

-debuginclusive debug_level

Specifies a debug level for controlling which debug messages get rendered in the expanded output. Since it's "inclusive" all messages with a level at the specified level and below will get rendered in the expanded output.

The following example expands the macro defined in the Macro section of the iApp, and sets the debug level to render all debug messages with a level of 11, 33 or 66:

```
tmsh::expand_macro -debug {11 33 66}
```

The following example expands the macro defined via a Tcl variable (mac), adds two variables (foo and enable_mything), and sets the debug level to render all messages of level 66 and below:

```
tmsh::expand_macro $mac -vars {foo bar enable_mything true} -debuginclusive 66
```

tmsh::get_ifile_text [iFile name]

Retrieve the text contained in the specified text iFile. When used on an iFile containing characters which are non-ascii or are not printable/space, an error will be returned.

tmsh::include [script name]

Runs the Tcl eval command on the specified script. The system evaluates the script at a global level, and all procedures in the included script are available to any other procedure. You must have previously created the script that is being included using the tmsh edit / cli script [name] command. If a full path is not given for the script name, tmsh will attempt to first locate the script from the same folder as the including script, then the root partition folder of the including script, and finally the /Common folder.

tmsh::run_proc [script_name:proc_name] options...

Runs the Tcl eval command on the specified script and process. The script script_name is loaded as if tmsh::include was called. After the script is loaded, the Tcl eval command is run on the specified Tcl

process. Any options that were specified are passed to the Tcl process. This is essentially a short form of running `tmsh::include script_name`, followed by running one of the Tcl processes contained in the script that was included.

The following example invokes the `display_pool_status` proc that is contained in the `pool_utils` script:

```
tmsh::run_proc pool_utils:display_pool_status
```

```
tmsh::stateless [disabled | enabled]
```

Modifies the behavior of `tmsh::create` and `tmsh::delete`.

When stateless mode is disabled, an attempt to create an object that already exists in the configuration results in an error, and an attempt to delete an object that does not exist in the configuration is an error.

When stateless mode is enabled, an attempt to create an object that already exists in the configuration does not result in an error, and an attempt to delete an object that does not exist in the configuration does not result in an error.

Enabling stateless mode enables scripts to successfully run multiple times with the same input.

The default value is disabled.

```
tmsh::version
```

Returns the version number of the BIG-IP system as a Tcl string. The version consists of three digits: a major, minor, and maintenance version, separated by periods. For example, 10.1.0 indicates minor version 1 of major version 10.

CONTEXT SENSITIVE HELP

Use the following commands to create a script that provides context sensitive help when a user types Tab or question mark (?).

```
script::help
```

Scripts can provide the `script::help` procedure. `tmsh` invokes the procedure when a user types a question mark (?) while entering the command sequence `run / cli script [name]`. If the specified script includes the `script::init` procedure, `tmsh` invokes it before the `script::help` procedure. The script can add context sensitive help by calling the `tmsh::add_help` and `tmsh::builtin_help` procedures. `tmsh` formats the help and displays it.

```
script::tabc
```

Scripts can provide the `script::tabc` procedure. The system invokes this procedure when the user types Tab while entering the command sequence `run / cli script`. If the `script::init` procedure is included in the script, that procedure is invoked before the `script::tabc` procedure. The script can add tab completion datasets to the script by calling the `tmsh::add_tabc` and `tmsh::builtin_tabc` procedures. `tmsh` either formats and displays the tab completion datasets, or if possible, completes the current argument.

```
tmsh::csh
```

`tmsh::csh` is a Tcl string variable that can be used in the `script::init` procedure to determine the context in which the `script::init` procedure was invoked.

`tmsh::csh` is set to one of the following:

question mark (?)
Indicates that the user typed a question mark (?).

TABC Indicates the user pressed the Tab key.

an empty string ""
Indicates the script is being run.

```
tmsh::add_help [ [category item description] | [description] ]
```

Displays context sensitive help when the user types a question mark (?). If you supply one argument, that argument displays as-is with no formatting applied to the description.

If you supply three arguments, one or more datasets are constructed. The first argument is the name of the dataset. The second argument is an item in the dataset. The third argument is a description of the item. This command has an effect only if the Tcl `tmsh::csh` variable is set to question mark (?).

```
tmsh::builtin_help ["tmsh command" args...]
```

Presents the same results as typing a question mark (?) while entering a `tmsh` command. The system stores a set of possible completions and displays the possibilities when the `script::help` procedure returns. This command has an effect only if the Tcl `tmsh::csh` variable is set to question mark (?).

```
tmsh::add_tabc [ [category item] | [item] ]
```

Adds tab completion datasets. If you supply one argument, the system adds that argument to an anonymous dataset. If you supply two arguments, the system constructs one or more datasets. The first argument is the name of the dataset. The second argument is an item in the dataset. Potential completions are displayed in groups based on category. This command has an effect only if the Tcl `$tmsh::csh` variable is set to TABC.

```
tmsh::builtin_tabc ["tmsh command" args...]
```

Many of the `tmsh` commands that are available for scripting are also available in the interactive shell. A script can use the `tmsh::builtin_tabc` command to present the same tab completion results as a built-in command. The command does not return a value. The set of possible completions are stored internally and displayed when the `script::tabc` procedure returns. This command has an effect only if the Tcl `$tmsh::csh` variable is set to TABC.

THIRD PARTY TCL LIBRARY USAGE

A selection of third party libraries have been tested to work within the CLI script environment, including MD5, BASE64, SHA1/SHA256, HTTP, TLS, TCL Perl, LDAP client, and XML parser. The TCL packages can only reside in the directory of /usr/share/compat-tcl8.4.

Important: Only these tested packages are supported currently.

This example demonstrates the use of a Tcl package command to make use of tls/https. The TLS package is installed in the directory /usr/share/compat-tcl8.4/tls in the form of two files: tls.tcl and libtls1.6.1.so.

Modify script /Common/use_tls {

```
proc script::run {} {
    set pkg_name tls
    set pkg_version 1.6
    package require http
    if {[catch {package require $pkg_name pkg_version}]} {
    puts "No package found: $pkg_name!\n"
    } else {
    puts "Found package: $pkg_name!\n"
    http::register https 443 tls::socket
    set token [http::geturl https://172.27.42.161/]
    upvar #0 $token state
    puts $state(http)
    puts $state(body)
    }
}
```

This example uses the callback function to handle http data.

```
cli script /Common/use_http2 {
proc script::httpCallback {token} {
    upvar #0 $token state
    puts $state(http)
    puts $state(body)
    incr ::got_something
}
proc script::run {} {
    namespace eval :: {
    set got_something 0
    }
    set pkg_name http
    set pkg_version 2.4.5
    if {[catch {package require $pkg_name $pkg_version}]} {
    puts "No package found: $pkg_name!\n"
    } else {
    puts "Found package: $pkg_name!\n"
    http::geturl http://172.27.42.22/index.htm -command script::httpCallback
    vwait ::got_something
    }
}
```

This example uses the LDAP client package to query data.

```
cli script /Common/use_ldap {
proc script::run {} {
    set pkg_name ldap
    if {[catch {package require $pkg_name 1.8}]} {
    puts "No package found: $pkg_name!\n"
    } else {
    puts "Found package: $pkg_name!\n"
    set handle [ldap::connect 172.27.1.2]
    ldap::bind $handle
    set results [ldap::search $handle "dc=f5,dc=com" "(uid=test)" {}]
    foreach result $results {
        puts $result
    }
    ldap::unbind $handle
    ldap::disconnect $handle
    }
}
```

Here are some additional examples:

```
cli script /Common/use_parray {
proc script::run {} {
    puts [info patch]
    namespace eval :: {
        set pkg_location /usr/share/compat-tcl8.4/
        source [file join $pkg_location package.tcl]
    }
    puts "NS: [namespace current]"
    set pkg_location $::pkg_location
    source [file join $pkg_location parray.tcl]
}
```

```

parray ::tcl_platform
}
}

cli script /Common/use_sha2 {
proc script::run {} {
set pkg_name sha256
if {[catch {package require $pkg_name}]} {
puts "No package found: $pkg_name!\n"
} else {
puts "Found package: $pkg_name!\n"
puts "TCL does SHA2 now:"
puts [sha2::sha256 "TCL does SHA2"]
}
}
}

cli script /Common/use_tclperl {
proc script::run {} {
set pkg_name tclperl
if {[catch {package require $pkg_name}]} {
puts "No package found: $pkg_name!\n"
} else {
puts "Found package: $pkg_name!\n"
set interpreter [perl::interp new]
$interpreter eval {print "Hello World\n"}
perl::interp delete $interpreter
}
}
}
}

```

SPECIAL CHARACTERS

There are several characters that are part of both Tcl and tmsh syntax. You must escape these characters in a shell script so that Tcl passes them to tmsh. You can use standard Tcl escape characters, such as quotes and back slashes. You must escape curly braces ({ }), for example, "{ " }".

```
tmsh::create ltm pool my_pool members add "{ 10.1.2.3:80 }"
```

Creates a Local Traffic Manager pool named my_pool.

DISABLED COMMANDS

The following commands are disabled for users that have not been assigned a user role of Administrator or Resource Administrator:

- Â· auto_execok
- Â· auto_import
- Â· auto_load
- Â· auto_mkindex
- Â· auto_mkindex_old
- Â· auto_qualify
- Â· auto_reset
- Â· bgerror
- Â· cd
- Â· close
- Â· eof
- Â· exec
- Â· fblocked
- Â· fconfigure
- Â· fcopy
- Â· file
- Â· fileevent
- Â· filename
- Â· flush
- Â· glob
- Â· http
- Â· interp

- Â· load
- Â· memory
- Â· open
- Â· package
- Â· pid
- Â· pkg:create
- Â· pkg_mkindex
- Â· pwd
- Â· seek
- Â· socket
- Â· source
- Â· tcl_findLibrary
- Â· tell
- Â· unknown
- Â· updates
- Â· vwait

EXAMPLES

The following example demonstrates the use of all tmsh Tcl commands. The script displays all configuration property values or all status and statistic values for the specified component, depending on the specified arguments. The system displays all configuration settings if you replace [tmsh::get_config \$comp all-properties] with [tmsh::get_config / all-properties]. The use of the all-properties option ensures that all options are displayed.

This command sequence is an example of how to run the following script: run / cli script example.tcl config ltm pool.

```
cli script example.tcl {
    proc script::init { } {
set ::field_fmt "%-25s %s"
set ::usage_string "usage: [lindex $tmsh::argv 0] \
"
    }

    proc script::help { } {
if { $tmsh::argc < 2 } {
tmsh::add_help Options: config "Display configuration"
tmsh::add_help Options: status \
"Display status and statistics"
}
else {
build_csh tmsh::builtin_help
}
    }

    proc script::tabc { } {
if { $tmsh::argc < 2 } {
tmsh::add_tabc config
tmsh::add_tabc status
}
else {
build_csh tmsh::builtin_tabc
}
    }

    proc script::run { } {
if { $tmsh::argc < 3 } {
usage
}
set opt [lindex $tmsh::argv 1]
if { $opt != "config" && $opt != "status" } {
usage
}
set comp ""
for {set idx 2} {$idx < $tmsh::argc} {incr idx} {
append comp "[lindex $tmsh::argv $idx] "
}

if { $opt == "config" } {
set objs [tmsh::get_config $comp all-properties]
}
}
```

```

else {
    set objs [tmsh::get_status $comp]
}

set idx 0
set total [llength $objs]

while { $idx < $total } {
    set obj [lindex $objs $idx]
    print_object obj
    puts ""
    incr idx;
}

proc print_fields { objVar } {
upvar $objVar obj
set fdx 0
set fields [tmsh::get_field_names value $obj]
set field_count [llength $fields]
while { $fdx < $field_count } {
    set field [lindex $fields $fdx]
    puts [format ::field_fmt $field \
[tmsh::get_field_value $obj $field]]
    incr fdx
}
}

proc print_object { objVar } {
upvar $objVar obj
puts "[tmsh::get_type $obj] [tmsh::get_name $obj]"

# name/value pairs
print_fields obj

# nested objects
set fdx 0
set fields [tmsh::get_field_names nested $obj]
set count [llength $fields]
while { $fdx < $count } {
    set field [lindex $fields $fdx]
    set nested_objects [tmsh::get_field_value $obj $field]
    set ndx 0
    set n_count [llength $nested_objects]
    while { $ndx < $n_count } {
set nobj [lindex $nested_objects $ndx]
print_object nobj
incr ndx
    }
    if { $n_count == 0 } {
puts [format ::field_fmt $field "none"]
    }
    incr fdx
}
}

proc build_csh { command } {
# generate context sensitive help, tab completion or "?"
set args ""
for {set idx 2} {$idx < $tmsh::argc} {incr idx} {
    lappend args [lindex $tmsh::argv $idx]
}
set opt [lindex $tmsh::argv 1]
if { $opt == "config" } {
    $command list $args
}
elseif { $opt == "status" } {
    $command show $args
}
else {
    puts "\nunexpected argument: $opt"
}
return $args
}

proc usage { } {
puts ::usage_string
exit
}
}

```

EXAMPLES

The following example demonstrates the scripts using 11.6.0 tmsh syntax.

```

cli script example_ver.tcl {
...
proc script::run {} {

```

```

tmsh::modify cli version active 11.6.0
# the tmsh command in the section below should contain 11.6.0 tmsh syntax
puts stdout [tmsh::list ltm pool mypool]
...
}
}

```

The following example demonstrates the scripts that are using multiple tmsh syntaxes.

```

cli script example_multi_ver.tcl {
....
proc script::run {} {
tmsh::modify cli version active 11.5.0
# the tmsh command in the section below should contain 11.5.0 tmsh syntax
puts stdout [tmsh::list ltm pool mypool]
...

tmsh::modify cli version active 11.6.0
# the tmsh command in the section below should contain 11.6.0 tmsh syntax
puts stdout [tmsh::list ltm pool mypool]
...
}
}

```

SEE ALSO

cli alias, create, delete, edit, glob, list, modify, regex, reset-stats, show, tmsh and generate.

For complete information about tmsh, see the Traffic Management Shell (tmsh) Reference Guide. This guide is available on the Ask F5(sm) Knowledge Base (www.askf5.com).

For information about Tcl, see www.tcl.tk.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2016. All rights reserved.

BIG-IP 2016-03-14 cli script(1)

cli transaction

NAME

transaction - Opens batch mode within which you can submit a set of commands as a single transaction.

MODULE

cli

SYNTAX

Use the transaction component within the cli module to open batch mode, enter a series of commands, and then submit the commands as a single transaction.

CREATE/MODIFY

create transaction

modify transaction

options:

delete [entry_id]

submit transaction [validate-only]

DISPLAY

list transaction

DELETE

delete transaction

DESCRIPTION

tmsh parses each command that you enter in batch mode. If the command passes a syntax check, tmsh saves it as part of the transaction you are creating and returns a confirmation. After you finish adding commands, you submit the transaction to change the running configuration of the system. You must run the save config command to save the changes to the stored configuration files.

If, while creating a transaction, you decide you do not want to change the running configuration, you can delete the transaction rather than submit it. However, you can recreate a transaction that you have deleted by using the cli history component.

There are a few commands that you can enter on the command line that the system immediately runs, rather than

adding the commands to a transaction. These commands are list and show. Additionally, tmsh immediately runs the command sequence run bigpipe, but does not add it to the transaction.

EXAMPLES

The following example shows the commands that you enter from within the ltm module to create and submit a transaction that creates a Local Traffic Manager pool and virtual server, and then associates the two.

1. Open tmsh batch mode:

```
create /cli transaction
```

2. Add a command to the transaction that creates pool1 for the Local Traffic Manager using the default values for a pool:

```
create pool pool1
```

3. Add a command to the transaction that creates the virtual server virtual1 for the Local Traffic Manager using the default values for a virtual server, and associates it with pool1.

```
create virtual virtual1 pool pool1
```

4. Display, in a numbered list, the current set of commands in the transaction:

```
list /cli transaction
```

Note: You can use the preceding command to determine the entry ID of a command. Then, you can use this ID to remove or replace a command in the transaction, or to identify a command before which you want to insert another command.

5. Submit the transaction:

```
submit /cli transaction
```

OPTIONS

command

Specifies, in quotation marks, the full path to a command to add to or delete from the transaction that you are creating. You can also replace an existing command with another command or insert a command before a command in the transaction.

create

Opens batch mode.

delete

Deletes the transaction that you are creating and closes batch mode.

list Displays, in a numbered list, the current set of commands in the transaction that you are creating.

modify

Specifies a previously entered line in the transaction that you want to change. The options are:

delete

Deletes the specified entries from the transaction that you are creating.

entry_id

Specifies the number of a command in the list of commands in the transaction that you want to delete.

submit [validate-only]

Submits the transaction that you are entering and closes batch mode. The transaction is submitted in the context of the cli admin-partitions settings that are active when the submit command is issued.

validate-only

Validates the configuration changes without putting them in production.

SEE ALSO

cli admin-partitions, create, delete, list, modify, submit, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2015-06-01 cli transaction(1)

cli version

NAME

version - Displays and Configures tmsh versions.

MODULE

cli

SYNTAX

Configure the version component within the cli module using the syntax shown in the following sections.

MODIFY

modify version [option]

options:

active [string]

DISPLAY

show version

DESCRIPTION

You can use the version component to configure tmsh to run the specified version.

EXAMPLES

modify cli version active 11.5.0

Configures tmsh run 11.5.0 version.

show cli version

Displays the latest, active and supported versions of TMSH.

OPTIONS

active

Specifies the active version of TMSH.

latest

Displays TMSH the latest version. This is used as the default version.

supported

Displays the current supported TMSH versions on the system.

imported

Displays the imported TMSH versions on the system. An imported TMSH version will be imported from a UCS created from TMSH version which is not supported in the current system - a very rare case. Be aware, for an imported TMSH version, only syntax is supported, if it requires other handling other than syntax change, it will not supported. So, for an imported TMSH version, it is not fully supported. By default, this entry will not be displayed unless preference is set.

SEE ALSO

show, modify, sys ucs, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-06-11 cli version(1)

cm

cm add-to-trust

NAME

add-to-trust - Add a device to a trust domain.

MODULE

cm

SYNTAX

Run the add-to-trust program within the cm module using the syntax in the following section. The trust-domain name 'Root' is optional beginning version 13.0.0.

MODIFY

run add-to-trust [Root]

options:

[ca-device | non-ca-device]

device [string]

port [port_number]

device-name [string]

password [string]

username [string]

DESCRIPTION

You can use the add-to-trust command to add a device to a trust domain. This is an alternate helper command to the modify trust-domain ca-devices|non-ca-devices add ... command.

EXAMPLES

```
run add-to-trust ca-device device 10.20.30.40 device-name peer1 username homer password illiad
```

Adds a device to the list of ca-devices in the trust domain.

OPTIONS

ca-device

Indicates that the added device is a certificate authority device.

device

Indicates the FQDN or the management-ip of the device being added to the trust domain

port Device port number if other than 443 when adding new device. This parameter is optional.

device-name

Used to specify the name of a new device.

md5-fingerprint

SSL certificate md5 fingerprint is deprecated beginning version 13.0.0. Use sha1-fingerprint.

non-ca-device

Indicates that the added device is a subordinate device. The target device cannot be used as a signing authority.

password

Specifies the password for a new device.

serial

SSL certificate serial number is deprecated beginning version 13.0.0. Use sha1-fingerprint.

sha1-fingerprint

Specifies the SSL certificate (DER format) sha1 fingerprint when verifying the identity of a new device. This field is optional.

signature

SSL certificate signature is deprecated beginning version 13.0.0. Use sha1-fingerprint.

username

Specifies the user name required to log on to a device when adding the device to the trust domain. The user "root" is invalid, and will be disallowed. Any user that has administrator privileges and can use iControl can be used here.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-10-04 cm add-to-trust(1)

cm cert

NAME

cert - Manages a CM trust certificate file.

MODULE

cm

SYNTAX

Display the cert component within the cm module using the syntax shown in the following sections.

DISPLAY

list cert

list cert [[[name] | [glob] | [regex]] ...]

show running-config cert

show running-config cert [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

certificate-key-size

checksum

create-time

created-by

email

expiration-date
expiration-string
fingerprint
is-bundle
issuer
key-type
last-update-time
mode
non-default-properties
one-line
partition
recursive
revision
serial-number
size
source-path
subject
subject-alternative-name
system-path
updated-by
version

DESCRIPTION

You can use the cert component to display CM trust certificates.

OPTIONS

app-service

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

certificate-key-size

Displays the number of bits in the key associated with this certificate.

checksum

Displays a cryptographic hash or checksum of the file contents for use in verification of file integrity.

create-time

Displays the time at which the trust certificate was created.

created-by

Displays the name of the person, who originally created the trust certificate.

email

Displays the email of the person, who originally created the trust certificate.

expiration-date

Displays the date at which the trust certificate expires. The date is stored as a POSIX time.

expiration-string

Displays a string representation of the trust certificate expiration date.

fingerprint

Specifies the cryptographic fingerprint of the trust certificate.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

is-bundle

Indicates whether the trust certificate file is a bundle (that is, whether it contains more than one certificate).

issuer

Displays the X.509 information for the issuer of the trust certificate. If the trust certificate is a bundle, then this displays the issuer information for the primary (first) trust certificate in the bundle.

key-type

Displays the type of cryptographic key associated with this trust certificate.

last-update-time

Displays the last time the trust certificate was modified.

mode Displays the UNIX(r) file permissions mode for the file associated with this trust certificate as a numerical value.

partition

Displays the partition within which the trust certificate file resides.

recursive

Displays all objects of the specified type and the folder that contains the object.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

revision

Displays the number of the latest revision of the trust certificate. The revision starts with 1 and increments on each update.

serial-number

Displays the serial number of the trust certificate.

size Displays the size (in bytes) of the file associated with the trust certificate.

source-path [URL]

Displays the path to the source of the trust certificate as a URL, for example:

source-path http://cert-server/cert_store/certs/vs_132.key

source-path https://cert-server/cert_store/certs/vs_132.key

source-path ftp://username:password@server/cert_store/certs/vs_132.key

subject

Displays X.509 information about the subject of the trust certificate. If the certificate is a bundle, then the subject information for the primary (first) trust certificate in the bundle displays.

subject-alternative-name

Displays a standard X.509 extension as shown in RFC 2459.

system-path

Displays the path to the trust certificate.

updated-by

Displays the name of the person, who last updated the trust certificate.

version

Displays the X.509 version of the trust certificate.

SEE ALSO

glob, list, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2015-03-17 cm cert(1)

cm config-sync

NAME

config-sync - Manually synchronizes the configuration between devices.

MODULE

cm

SYNTAX

Run the config-sync program within the cm module using the syntax in the following section.

MODIFY

run config-sync

options:

from-group

recover-sync

to-group

force-full-load-push

DESCRIPTION

This command starts a configuration synchronization job. One of the from-group, to-group, or recover-sync options must be used to specify the direction of the synchronization.

EXAMPLES

run config-sync from-group my_dg

Updates the configuration on the local device with the configuration from the remote device in the device group /Common/my_dg with the newest configuration. If the local device already has the newest configuration, then the configuration synchronization does nothing.

run config-sync to-group my_dg

Updates the configurations on the remote devices in the device group /Common/my_dg with the configuration on

the local device. If the local device does not have the newest configuration, then the configuration synchronization does nothing.

run config-sync recover-sync

Resets the local device configuration and restores the trust domain, device, and device-group information to default settings.

OPTIONS

from-group

Updates the configuration of the local device with the configuration of the remote device in the specified device group that has the newest configuration. If the local device already has the newest configuration, then the configuration synchronization does nothing. This option is mutually exclusive of the to-group and recover-sync options.

recover-sync

WARNING: If you have any local-only configuration, do not use this option.

This option will delete all configuration except for that which is necessary to remain in your device groups. At this point, you can sync from those groups to this device in order to restore your configuration.

Resets the local device configuration and restores the trust domain, device, and device-group information to default settings. After this recovery, you can sync the local device with its peers by running config-sync on a peer device and specifying the device group in which the local device is a member. This option is mutually exclusive of the from-group and to-group options.

to-group

Updates the configurations of the remote devices in the specified device group with the configuration of the local device. If the local device does not have the newest configuration, then the configuration synchronization does nothing. This option is mutually exclusive of the from-group and recover-sync options.

force-full-load-push

This option may only be used if to-group has also been specified. It forces all other devices to pull all synchronizable configuration from this device, even if those devices have the same or newer configuration. It will always send the entire configuration, rather than just the changes, regardless of the device group's full-load-on-sync setting.

You may lose configuration by using this option. This happens if you force a push even when another device in the device group has newer configuration.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2011, 2013, 2015. All rights reserved.

BIG-IP 2017-10-09 cm config-sync(1)

cm device-group

NAME

device-group - Configures device groups. The device groups that are created and maintained by the system cannot be modified by the user.

MODULE

cm

SYNTAX

Modify the device-group component within the cm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create device-group [name]
```

```
modify device-group [name]
```

options:

```
app-service [[string] | none]
```

```
asm-sync [ enabled | disabled ]
```

```
auto-sync [ enabled | disabled ]
```

```
description [string]
```

```
devices [add | delete | modify | replace-all-with] {  
  [ device_name ]
```

```
}
```

```
full-load-on-sync [true | false]
```

```
incremental-config-sync-size-max [integer]
```

```
network-failover [ enabled | disabled ]
save-on-auto-sync [ true | false ]
type [ sync-only | sync-failover ]
clear-incremental-config-sync-cache
```

```
edit device-group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list device-group
list device-group [ [ [name] | [glob] | [regex] ] ... ]
show running-config device-group
show running-config device-group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  partition
  recursive
```

```
show device-group
show device-group [name]
options:
  field-fmt
```

DELETE

```
delete device-group [name]
```

Note: The device group must be empty, and you must remove all references to the device group, before you can delete the device group.

DESCRIPTION

You can use the device-group component to manage sets of devices used for configuration synchronization and failover.

EXAMPLES

```
create device-group my_device_group devices add {
  /Common/device1
  /Common/device2
}
```

Creates a sync-only device group named my_device_group with two devices, device1 and device2.

```
delete device-group my_device_group
```

Deletes the device group named my_device_group.

```
list device-group my_device_group
```

Displays properties of the device group named my_device_group.

```
modify device-group my_device_group clear-incremental-config-sync-cache
```

Warning: Do not use this option without assistance from the F5 Technical Support team.

Clears the incremental configuration synchronization cache. The next configuration synchronization for my_device_group that pulls configuration from this device will be a full load.

OPTIONS

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

asm-sync

Specifies whether to synchronize ASM configurations of device group members. The default value is disabled. A device can be a member of only one ASM-enabled device group.

auto-sync

Specifies whether the device group automatically synchronizes configuration data to its members. The default value is disabled. Configuration will be saved on remote devices after receiving configuration updates if save-on-auto-sync is enabled.

clear-incremental-config-sync-cache

Warning: Do not use this option without assistance from the F5 Technical Support team.

The incremental configuration synchronization mechanism keeps a cache of transactions in each device group. Specifying this option will remove all transactions from the cache for the given device groups. This will not remove configuration from the device group, but will cause the next load in that group from the current device to be a full load.

description

Specifies a user-defined description of the device group.

devices

Adds, deletes, or replaces a set of devices to a device group by specifying the device name(s). When the local device is removed from a device group then all of the sys folders that are associated with the device group are reset to have no device group and the name of each folder that was updated is logged to /var/log/itm.

full-load-on-sync

Specifies that the entire configuration for a device group is sent when configuration synchronization is performed. The default value is false.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

incremental-config-sync-size-max

Specifies the maximum size (in KB) to devote to incremental config sync cached transactions. The default is 1024 KB. Valid range is between 128 and 10240."

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

network-failover

When the device group type is failover, specifies whether network failover is used.

partition

Displays the administrative partition within which the device group resides.

recursive

Displays all objects of the specified type and the folder that contains the object.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

save-on-auto-sync

Specifies whether to save the configuration on the remote devices following an automatic configuration synchronization. A device group configured for manual synchronization will always save on the remote devices regardless of this setting.

type Specifies the type of device group. You can use this option only when you create a device group. You cannot modify the type of a device group. The default value is sync-only.

SEE ALSO

create, delete, device, edit, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2017-05-25 cm device-group(1)

cm device

NAME

device - Manages a device.

MODULE

cm

SYNTAX

Manage the device component within the cm module using the syntax shown in the following sections. The 'create cm device' and 'delete cm device' commands are deprecated beginning version 13.0.0.

MODIFY

modify device [name]

options:

comment [string]

configsnc-ip [ip address | none]

contact [string]

description [string]

ha-capacity [integer]

location [string]

mgmt-unicast-mode [both | ipv4 | ipv6]

mirror-ip [ip address | any6]

mirror-secondary-ip [ip address | any6]

multicast-interface [string]

multicast-ip [ip address]

```
multicast-port [integer]
unicast-address [ none | {
  {
    ip [ip address]
    port [port number]
    effective-ip [ip address]
    effective-port [port number]
  }
  ...
}
]
```

```
edit device [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list device
```

```
list device [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config device
```

```
show running-config device [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  active-modules
  all-properties
  app-service
  base-mac
  build
  cert
  chassis-id
  chassis-type
  failover-stats
  hostname
  inactive-modules
  key
  location
  management-ip
  marketing-name
  non-default-properties
  one-line
  optional-modules
  partition
  platform-id
  product
  recursive
  self-device
  time-limited-modules
  time-zone
  version
```

```
show device-group
```

```
show device-group [name]
```

```
options:
  all
  field-fmt
```

DESCRIPTION

You can use the device component to manage devices.

Warning: To add or remove devices on the BIG-IP system, modify the Root trust domain. For more information, see help trust-domain.

OPTIONS

active-modules
Displays the licensed modules that are currently active on the device.

app-service
Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify the object. Only the Application Service can modify the object.

base-mac
Displays the base MAC address for the device.

build
Displays the software build number.

cert Displays the identity certificate used for device trust.

chassis-id
Displays the chassis identifier.

chassis-type
Displays the chassis type. The possible values are individual and viprion.

comment
Specifies user comments about the device.

`configs-sync-ip`

Specifies the IP address used for configuration synchronization. If you specify a self IP address, the self IP address object must be located in the Common folder.

`contact`

Specifies administrator contact information.

`description`

Specifies a user-defined description of the device.

`edition`

Displays the software edition.

`failover-state`

Displays the device failover state.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`ha-capacity`

Specifies a number that represents the relative capacity of the device to be active for a number of traffic groups. This value along with the traffic group's ha-load-factor is used by the failover daemon to make traffic groups active amongst the available devices. The value is zero by default which means the device may run any number of traffic groups. The value must be within a valid range: 0 - 100000 inclusive.

`hostname`

Displays the hostname of the device.

`inactive-modules`

Displays the licensed modules that are currently inactive on the device.

`key` Displays the identity key used for device trust.

`location`

Specifies the physical location of the device.

`marketing-name`

Displays the marketing name of the device platform.

`mgmt-unicast-mode`

When a management-ip is added as a failover unicast address; this attribute allows to specify using only an specific family instead of the default both.

`mirror-ip`

Specifies the IP address used for state mirroring. If you specify a self IP address, the self IP address object must be located in the Common folder.

`mirror-secondary-ip`

Specifies the secondary IP address used for state mirroring. If you specify a self IP address, the self IP address object must be located in the Common folder.

`multicast-interface`

Specifies the interface name used for the failover multicast IP address.

`multicast-ip`

Specifies the multicast IP address used for failover.

`multicast-port`

Specifies the multicast port used for failover.

`optional-modules`

Displays the modules that are available for the current platform, but are not currently licensed.

`platform-id`

Displays the device platform identifier.

`product`

Displays the software product name.

`recursive`

Displays all objects of the specified type and the folder that contains the object.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`self-device`

Displays true, when the device is the self device.

`time-limited-modules`

Displays the licensed modules that are time-limited.

`time-zone`

Displays the time zone configured on the device.

unicast-address

If present, specifies the entire set of unicast addresses used for failover, and replaces all previous unicast addresses. The keyword none may be used to remove all unicast addresses.

Multiple unicast addresses may be specified, each of which has an ip, port, effective-ip, and effective-port.

Each unicast-address when specified has the following parameters:

ip The IP address that the failover daemon will listen on for packets from its peers. This address must be a non-floating self-IP or a management address.

The keyword management-ip may be used to specify that the current addresses of the device management interface are used. The mgmt-unicast-mode will narrow down the selection based on the address family.

port The IP port that the failover daemon uses to accept packets from its peers. If not specified, 1026 will be used.

effective-ip

The IP address that peers can use to reach this unicast address IP. This option is only needed if a address-translating firewall exists between the peer BIG-IPs. If not present, the effective-ip is the same as the ip.

effective-port

The port that peers can use to reach this unicast address. This option is only needed if a port-translating firewall exists between the peer BIG-IPs. If not present, the effective-port is the same as the port.

version

Displays the software version number.

SEE ALSO

edit, glob, list, modify, mv, regex, show, tmsh, trust-domain

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-10-25 cm device(1)

cm failover-status

NAME

failover-status - Display the failover status of the local device.

MODULE

cm

SYNTAX

Display failover-status component within the cm module using the syntax in the following section.

DISPLAY

show failover-status

options:

field-fmt

DESCRIPTION

You can use the failover-status component to display the failover status of the local device.

For information about the options that you can use with the command show, see help show.

EXAMPLE

show failover-status

Displays the failover status of the local device.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

cm key

NAME

key - Manages a CM trust certificate private key file.

MODULE

cm

SYNTAX

Display or delete a key component within the cm module using the syntax shown in the following sections.

DISPLAY

list key

list key [[[name] | [glob] | [regex]] ...]

show running-config key

show running-config key [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

checksum

create-time

created-by

key-size

key-type

last-update-time

mode

non-default-properties

one-line

partition

recursive

revision

security-type

size

source-path

system-path

updated-by

DELETE

delete key [name]

DESCRIPTION

You can use the following options with the key component.

OPTIONS

app-service

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

checksum

Displays a cryptographic hash or checksum of the key for use in verification of key integrity.

create-time

Displays the time at which the key was created.

created-by

Displays the user who originally created the key.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

key-size

Displays the size of the cryptographic key, in bits.

key-type

Displays the cryptographic algorithm that this key is compatible with. A key can be one of two types:

rsa-private

The key is an RSA private key.

dsa-private

The key is a DSA based private key.

last-update-time

Displays the time at which the key was last modified.

mode Displays the UNIX file permissions mode for the file associated with this key. The mode is expressed in numerical form.

name Specifies the name of the key you want to delete.

partition
Displays the partition within which the key resides.

recursive
Displays all objects of the specified type and the folder that contains the object.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

revision
Displays the latest revision of the key. The revision starts with 1 increments on each update.

security-type
Displays the type of security used to handle or store the key. There are four mutually exclusive options:

normal
Indicate the key resides in a standard form on the file-system. This is the default security type.

fips Indicates that the key is protected by a FIPS device on the system, and is only applicable to devices with FIPS support.

password
Indicates that the key is protected by a passphrase and stored in encrypted form.

nethsm
Indicates that the key is protected by a FIPS device outside the system.

size Displays the size (in bytes) of the file associated with this file object.

source-path
Displays the location (URI) from where the file will be copied.

system-path
Displays the location where the key is stored on the system.

updated-by
Displays the name of the user who last updated the key.

SEE ALSO

delete, glob, list, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2018-03-29 cm key(1)

cm remove-from-trust

NAME

remove-from-trust - Remove a device from a trust domain.

MODULE

cm

SYNTAX

Run the remove-from-trust program within the cm module using the syntax in the following section. The trust-domain name 'Root' is optional beginning version 13.0.0.

MODIFY

```
run remove-from-trust [Root]
options:
[ca-device | non-ca-device]
device-name [string]
```

DESCRIPTION

You can use the remove-from-trust command to remove a device from a trust domain. This is an alternate helper command to the modify trust-domain ca-devices|non-ca-devices delete ... command.

EXAMPLES

```
run remove-from-trust ca-device device-name peer1
```

Removes a device from the list of ca-devices in the trust domain Root.

OPTIONS

ca-device

Indicates that the added device is a certificate authority device.

device-name

Used to specify the name of a new device.

non-ca-device

Indicates that the added device is a subordinate device.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014-2015. All rights reserved.

BIG-IP 2016-02-17 cm remove-from-trust(1)

cm sha1-fingerprint

NAME

sha1-fingerprint - Display SHA1 fingerprint of the local device SSL certificate in DER format.

MODULE

cm

SYNTAX

Run the sha1-fingerprint command sequence within the cm module using the syntax in the following section.

DISPLAY

```
show sha1-fingerprint
```

options:

field-fmt

DESCRIPTION

You can use the sha1-fingerprint component to display the local box SSL SHA1 checksum and provide this value to a remote authority device that is trying to add the local device to a trust domain. This allows the remote device to verify the identity of the local device when establishing the SSL connection.

For information about the options that you can use with the command show, run the command sequence help show.

EXAMPLE

```
show sha1-fingerprint
```

Display SHA1 fingerprint of the local device SSL certificate in DER format.

OPTIONS

field-fmt

Formats the sha1 fingerprint output in command syntax.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-01-20 cm sha1-fingerprint(1)

cm sniff-updates

NAME

sniff-updates - Displays the commit ID updates that occur over the CMI communications channel

MODULE

cm

SYNTAX

```
run cm sniff-updates
options:
[-v]
```

DESCRIPTION

You can use the sniff-updates program to monitor the internal CMI communications channel for commit ID updates. The system displays each update as it arrives, one per line.

```
(1) (2) (3) (4) (5) (6) (7) (8) (9) (10)
[15:35:57] bigip1 (v0.0.0) -> device_trust_group: CID 105.105 (bigip2) at 15:34:39 FORCE_SYNC
```

Output fields:

- 1) Time that update arrived from network
- 2) Source device
- 3) Version of source device
- 4) Destination devicegroup
- 5) CommitId ID
- 6) DeviceData CommitId ID
- 7) CommitId originator
- 8) CommitId timestamp
- 9) FORCE_SYNC if set (nothing if not)
- 10) Last sync error message (nothing if last sync was successful)

OPTIONS

You can use the following option when you run the sniff-updates program:

-v Formats the update output using fully-qualified device and device group names and exact time64_t timestamps.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 cm sniff-updates(1)

cm sync-status

NAME

sync-status - Displays the configuration synchronization status of the local device.

MODULE

cm

SYNTAX

Run the sync-status command sequence within the cm module using the syntax in the following section.

```
DISPLAY
show sync-status
options:
field-fmt
```

DESCRIPTION

You can use the sync-status component to display the configuration synchronization status of the local device.

For information about the options that you can use with the command show, run the command sequence help show.

EXAMPLE

```
show sync-status
```

Displays the configuration synchronization status of the local device:

OPTIONS

field-fmt

Formats the status output in command syntax.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-03-19 cm sync-status(1)

cm traffic-group

NAME

traffic-group - Manages a CM traffic group.

MODULE

cm

SYNTAX

Manage the traffic-group component within the cm module using the syntax shown in the following sections.

CREATE/MODIFY

create traffic-group [name]

modify traffic-group [name]

options:

app-service [[string] | none]

auto-failback-enabled [enabled | disabled]

auto-failback-time [integer]

description [string]

failover-method [ha-score | ha-order]

ha-group [string]

ha-load-factor [integer]

ha-order [string ...]

mac [mac address]

monitor { ha-group [string] }

edit traffic-group [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list traffic-group

list traffic-group [[[name] | [glob] | [regex]] ...]

show running-config traffic-group

show running-config traffic-group [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

is-floating

non-default-properties

one-line

partition

recursive

unit-id [integer]

show traffic-group

show traffic-group [name]

options:

all-properties

details

failover-objects

field-fmt

DELETE

delete traffic-group [name]

DESCRIPTION

You can use the traffic-group component to specify the failover behavior for devices in a failover device group.

EXAMPLES

create traffic-group my_traffic_group

Creates a traffic group named `my_traffic_group`.

```
create traffic-group my_traffic_group ha-order { my_device }
```

Creates a traffic group named `my_traffic_group` with a preferred device named `my_device`.

OPTIONS

`app-service`

Specifies the application service to which the object belongs. The default value is none. Note: If the `strict-updates` option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

`auto-failback-enabled`

Specifies whether the traffic group fails back to the default device.

`auto-failback-time`

Specifies the time required to fail back. The value must be within a valid range: 0 - 300 inclusive.

`details`

Only usable with the `show` command. Displays the active or next-active devices for this traffic group.

`failover-method`

Specifies the method used to decide if the current device needs to failover the traffic-group to another device. If the `failover-method` is set to `ha-score`, a score is calculated for each device in the monitoring HA group; and the highest score is the current active device while the next highest score will be the device that takes over if the current one fails. If the `failover-method` is set to `ha-order`, a list of devices and their respective HA load is used to decide the next one to take over if the current devices fails.

`failover-objects`

Only usable with the `show` command. Tells it to display all of the objects associated with that traffic group.

`ha-group`

Deprecated since v13.0.0. Use `monitor` and `failover-method` instead. This specifies a ha-group for the traffic group to decide the active device within the traffic group. The HA group must exist first. Note: This attribute is only specific to the local device i.e. not sync'ed to its peers in the traffic group. If you use this deprecated command attribute, it will also set the `failover-method` to `ha-score`.

`ha-order`

This list of devices specifies the order in which the devices will become active for the traffic group when a failure occurs. This list may contain zero, one or more entries up to the number of devices in the failover device group. If `auto-failback enabled` is set to true, this list must contain at least one entry for the auto-failback device.

`ha-load-factor`

Specifies a number for this traffic group that represents the load this traffic group presents to the system relative to other traffic groups. This allows the failover daemon to load balance the active traffic groups amongst the devices. The value is one by default. The value must be within a valid range: 1 - 1000 inclusive.

`monitor { ha-group [string] }`

This specifies a ha-group monitor for the traffic group to decide the active device within the traffic group. The HA group must exist first. Note: This attribute is only specific to the local device i.e. not sync'ed to its peers in the traffic group.

`description`

Specifies a user-defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`is-floating`

Indicates whether the traffic group can fail over to other devices in the device group.

`mac` Specifies a MAC address for the traffic group.

`partition`

Displays the administrative partition within which the device group resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`recursive`

Displays all objects of the specified type and the folder that contains the object.

`unit_id`

Displays the unit ID for the traffic group. The unit ID is set automatically when you create a traffic group. The value is between 1 and 15.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-10-04 cm traffic-group(1)

cm trust-domain

NAME

trust-domain - Manages a CM trust domain by providing control of object failover.

MODULE

cm

SYNTAX

Manage the trust-domain component within the cm module using the syntax shown in the following sections. The trust-domain name 'Root' is optional beginning with version 13.0.0. The 'create cm trust-domain' command is not allowed beginning version 13.0.0.

MODIFY

modify trust-domain [Root]

options:

```
add-device {
  [ ca-device [true | false] | non-ca-device [true | false] ]
  device-ip [string]
  device-port [port_number]
  device-name [string]
  username [string]
  password [string]
  sha1-fingerprint [string]
}
devices delete {
  [ device names ]
}
remove-device [string]
```

deprecated since v13.0.0:

```
ca-devices [add | delete | modify | replace-all-with] {
  [ device_name | ip address ]
}
md5-fingerprint [string]
name [string]
non-ca-devices [add | delete | modify | replace-all-with] {
  [ device_name | ip address ]
}
password [string]
serial [string]
sha1-fingerprint [string]
username [string]
```

DISPLAY

list trust-domain

list trust-domain [[[name] | [glob] | [regex]] ...]

show running-config trust-domain

show running-config trust-domain [[[name] | [glob] | [regex]] ...]

options:

```
all-properties
app-service
ca-cert
ca-cert-bundle
ca-key
ca-devices
non-ca-devices
non-default-properties
one-line
partition
recursive
status
trust-group
```

DELETE/RESTART

restart trust-domain

delete trust-domain

options:

```
keep-current-certificate-authority
import-user-defined-cert [string]
import-user-defined-key [string]
```

DESCRIPTION

You can use the trust-domain component to manage the behavior of objects during fail over.

DELETE/RESTART operations

When applied to a trust-domain these operations reset the trust and make this device standalone.

EXAMPLES

Adds a certificate authority:

```
modify trust-domain add-device { ca-device true device-ip 192.168.1.245 device-name myDevice1 device-port 1234
username admin password admin }
```

Adds a non-authoritative certificate:

```
modify trust-domain add-device { ca-device false ip 192.168.1.248 device-name myDevice2 username admin
password admin sha1-fingerprint ab7012e8d834e639f497b2b1c9f1e855a4dbe232 }
```

Removes a device from the trust domain:

```
modify trust-domain devices delete { myDevice1 myDevice2 }
```

Resets the trust and makes this device standalone:

```
restart cm trust-domain
```

or

```
delete cm trust-domain
```

EXAMPLES (deprecated since V13.0.0)

Adds a certificate authority:

```
modify trust-domain Root ca-devices add { 192.168.1.245 } name myDevice1 username admin password admin
```

Adds a non-authoritative certificate:

```
modify trust-domain Root non-ca-devices add { 192.168.1.245 } name myDevice1 username admin password admin
```

Removes a device from the trust domain:

```
modify trust-domain Root ca-devices delete { myDevice1 }
```

OPTIONS

add-device

Adds a device to the trust domain.

device-ip

Device IP address when adding new device.

device-port

Device port number if other than 443 when adding new device. This parameter is optional.

device-name

Device name when adding new device.

username

Specifies the user name required to log on to a device when adding the device to the trust domain.

password

Specifies the password corresponding to the username required to log on to a device when adding the device to the trust domain.

sha1-fingerprint

Specifies the SSL certificate (DER format) sha1 fingerprint when verifying the identity of a new device. This field is optional.

app-service

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

ca-cert

Displays the certificate authority device trust certificate.

ca-cert-bundle

Displays the bundled certificate authority device trust certificates used to authenticate incoming connections.

ca-devices

Create and modify operations are deprecated since v13.0.0. List operation is still supported. Set of certificate authority devices in the trust domain.

ca-key

Displays the certificate authority device trust key. This key only displays for certificate authorities.

devices

Removes one or more devices from the trust domain. It takes the name of the device as the identifier.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`import-user-defined-cert`

Specifies the certificate to import and use as the trust domain's new certificate authority. You must also use the `import-user-defined-key` option to specify the corresponding key. This option cannot be specified alongside `keep-current-certificate-authority`.

`import-user-defined-key`

Specifies the key to import and use as the trust domain's new certificate authority key. You must also use the `import-user-defined-cert` option to specify the corresponding certificate. This option cannot be specified alongside `keep-current-certificate-authority`.

`keep-current-certificate-authority`

By default, resetting trust will generate a new certificate authority. Adding this option to the `delete` command will instead keep the current certificate authority. This option cannot be specified alongside `import-user-defined-cert` or `import-user-defined-key`.

`md5-fingerprint`

Deprecated since v13.0.0. Specifies the SSL certificate fingerprint when verifying the identity of a new device.

`name` Deprecated since v13.0.0. Option used to specify the name of a new device.

`non-ca-devices`

Create and modify operations are deprecated since v13.0.0. List operation is still supported. Set of subordinate devices in the trust domain.

`password`

Deprecated since v13.0.0. Specifies the password for a new device.

`recursive`

Displays all objects of the specified type and the folder that contains the object.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an `@` sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`remove-device`

Removes a single device from the trust domain. It takes the name of the device as the identifier.

`serial`

Deprecated since v13.0.0. Specifies the SSL certificate serial number when verifying the identity of a new device.

`sha1-fingerprint`

Deprecated since v13.0.0. Specifies the SSL certificate fingerprint when verifying the identity of a new device.

`signature`

Deprecated since v13.0.0. Specifies the SSL certificate signature, when verifying the identity of a new device.

`status`

Displays the status of the trust domain.

`trust-group`

Displays the device group associated with the trust domain.

`username`

Deprecated since v13.0.0. Specifies the user name required to log on to a device when adding the device to the trust domain.

SEE ALSO

`delete`, `edit`, `glob`, `list`, `modify`, `regex`, `restart`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-10-11 cm trust-domain(1)

NAME

watch-devicegroup-device - Displays information about the devices in the device group to which the local device belongs.

MODULE

cm

SYNTAX

Run the watch-devicegroup-device program within the cm module using the syntax shown in the following sections.

RUN

run watch-devicegroup-device

DESCRIPTION

You can use the watch-devicegroup-device program to view dynamic information about the synchronization of the devices in the device group to which the local device belongs. You can use this information to monitor or troubleshoot the devices.

By default, multiple devices with identical information are collapsed into a single row that displays in green. The devices column identifies the devices by the suffix of the configuration synchronization IP address configured on the device. For example, if the devices in a device group have the IP addresses 10.0.0.15 and 10.0.0.16, the IDs in this column will be 15 and 16. Use the c (collapse) command to deactivate/activate this behavior.

For example, when you make a change to a device, the change is identified by a commit ID (cid.id) that displays when you run the watch-devicegroup-device program.

Within the program, you can use the following keys:

Press h to see a list of available commands.

Press the back tick key (`) to exit the help page.

Press c to toggle the view from a collapsed view to a full view. The command gathers information from every device in the trust group. When all devices in the trust group report the same information the view is collapsed and one line, highlighted in green, displays the information. The devices included in the line are shown in the devices column. You can press c to see the full view, which displays each device on a separate line.

Press Ctrl-C to exit the program.

Press the arrow keys to navigate across the columns or down the rows.

The content in the columns includes:

devices

Displays the suffix of the configuration synchronization IP address configured on the device. For example, if the devices in a device group have the config-sync IP addresses 10.0.0.15 and 10.0.0.16, the IDs in this column will be 15 and 16.

devgroup

Displays the name of the device group to which the device belongs. Note: This can be a sync-only, failover, or trust device group.

device

Displays the device object name.

cid.id

Displays the commit ID, which is a configuration change identifier.

cid-orig

Displays the name of the device on which the configuration change was made.

cid.time

Displays the time the configuration change was made.

last_sync

Displays the time the device configuration was last synchronized with the device group.

The devices in the to-group of a configuration synchronization display the same time in this column. The local device that pushes the configuration to the other devices in the device group (to-group) has a different value in this column.

The devices in the from-group of a configuration synchronization display the same time in this column. The local device that receives the configuration from the other devices has a different value in this column. You can use this information to determine a rollback strategy.

SEE ALSO

run, tmsd, watch-sys-device, watch-trafficgroup-device

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

cm watch-sys-device

NAME

watch-sys-device - Displays information about the local device.

MODULE

cm

SYNTAX

Run the watch-sys-device program within the cm module using the syntax shown in the following sections.

RUN

run watch-sys-device

DESCRIPTION

You can use the watch-sys-device program to view dynamic information about the local device.

By default, multiple devices with identical information are collapsed into a single row that displays in green. The devices column identifies the devices by the suffix of the configuration synchronization IP address configured on the device. For example, if the devices in a device group have the IP addresses 10.0.0.15 and 10.0.0.16, the IDs in this column will be 15 and 16. Use the c (collapse) command to deactivate/activate this behavior.

Within the program, you can use the following keys:

Press h to see a list of available commands.

Press the back tick key (`) to exit the help page.

Press c to toggle the view from a collapsed view to a full view. The command gathers information from every device in the trust group. When all devices in the trust group report the same information the view is collapsed and one line, highlighted in green, displays the information. The devices included in the line are shown in the devices column. You can press c to see the full view, which displays each device on a separate line.

Press Ctrl-C to exit the program.

Press the arrow keys to navigate across the columns or down the rows.

The content in the columns includes:

devices

Displays the suffix of the configuration synchronization IP address configured on the device. For example, if the devices in a device group have the IP addresses 10.0.0.15 and 10.0.0.16, the IDs in this column will be 15 and 16.

name Displays the device object name.

platform

Displays the device platform.

build

Displays the software build installed on the device.

failover_state

Displays the high availability state (active or standby) of the device.

mgmt_ip

Displays the IP address of the management port on the device.

configsnc_ip

Displays the IP address on the device that is used for configuration synchronization.

unicast_ip

Displays the unicast IP address of the device.

multicast_ip

Displays the multicast IP address of the device.

mirror_ip

Displays the IP address used for configuration mirroring for the device.

mirror_secondary_ip

Displays the secondary IP address used for configuration mirroring for the device.

desc Displays a description of the device.

SEE ALSO

run, tmsh, watch-devicegroup-device, watch-trafficgroup-device

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-06-30 cm watch-sys-device(1)

cm watch-trafficgroup-device

NAME

watch-trafficgroup-device - Displays information about the traffic groups associated with devices in a device group.

MODULE

cm

SYNTAX

Run the watch-trafficgroup-device program within the cm module using the syntax shown in the following sections.

RUN

run watch-trafficgroup-device

DESCRIPTION

You can use the watch-trafficgroup-device program to view dynamic information about the failover status of the devices in a device group to which the local device belongs. You can use this information to monitor or troubleshoot the devices in the device group.

By default, multiple devices with identical information are collapsed into a single row that displays in green. The devices column identifies the devices by the suffix of the configuration synchronization IP address configured on the device. For example, if the devices in a device group have the IP addresses 10.0.0.15 and 10.0.0.16, the IDs in this column will be 15 and 16. Use the c (collapse) command to deactivate/activate this behavior.

Within the program, you can use the following keys:

Press h to see a list of available commands.

Press the back tick key (`) to exit the help page.

Press c to toggle the view from a collapsed view to a full view. The command gathers information from every device in the device group. When all devices in the device group report the same information the view is collapsed and one line, highlighted in green, displays the information. The devices included in the line are shown in the devices column. You can press c to see the full view, which displays each device on a separate line.

Press Ctrl-C to exit the program.

Press the arrow keys to navigate across the columns or down the rows.

The content in the columns includes:

devices

Displays the suffix of the configuration synchronization IP address configured on the device. For example, if the devices in a device group have the IP addresses 10.0.0.15 and 10.0.0.16, the IDs in this column will be 15 and 16.

traffic_group

Displays the name of the traffic group associated with the device.

device_name

Displays the device object name.

failover_state

Displays the high availability state (active or standby) of the device.

next_active

Displays True for the device that becomes active if the active traffic group fails over.

score

Displays a system-generated high availability score used to select the next active device.

SEE ALSO

run, tmsh, watch-sys-device, watch-devicegroup-device

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2017-07-20 cm watch-trafficgroup-device(1)

gtm

gtm datacenter

NAME

datacenter - Configures a Global Traffic Manager(tm) data center.

MODULE

gtm

SYNTAX

Configure the datacenter component within the gtm module using the syntax in the following sections.

CREATE/MODIFY

create datacenter [name]

modify datacenter [name]

options:

app-service [[string] | none]

contact [[name] | none]

description [string]

[disabled | enabled]

location [none | [physical location]]

prober-fallback [any-available | inside-datacenter | outside-datacenter | pool | none]

prober-pool [none | name]

prober-preference [inside-datacenter | outside-datacenter | pool]

metadata

[add | delete | modify] {

[metadata_name ...] {

value ["value content"]

persist [true | false]

}

}

edit datacenter [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

reset-stats datacenter

reset-stats datacenter [[[name] | [glob] | [regex]] ...]

DISPLAY

list datacenter

list datacenter [[[name] | [glob] | [regex]] ...]

show running-config datacenter

show running-config datacenter [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

show datacenter

show datacenter [name]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

detail

field-fmt

DELETE

delete datacenter [name]

DESCRIPTION

You can use the datacenter component to create, modify, display, or delete a data center.

EXAMPLES

create datacenter DC1

Creates a data center named DC1 with options set to the default values.

list datacenter DC1 all-properties

Displays all properties of the data center named DC1.

OPTIONS

app-service

Specifies the name of the application service to which the data center belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the data center. Only the application service can modify or delete the data center.

contact

Specifies the name of the administrator or the name of the department that manages the data center. The default value is none.

description

User defined description.

[disabled | enabled]

Specifies whether the data center and its resources are available for load balancing. The default value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

location

Specifies the physical location of the data center. The default value is none.

metadata

Specifies user-defined data to associate with a server. By default the persist attribute is set to true. This means the data is saved into the configuration file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

prober-fallback

Specifies the type of prober to use to monitor servers defined in this data center when the preferred type is not available. The default value is any-available.

prober-pool

Specifies a prober pool to use to monitor servers defined in this data center when either the prober-preference or prober-fallback value is pool. If neither the prober-preference or prober-fallback value is pool, the prober-preference and prober-fallback values are set to pool and any-available.

prober-preference

Specifies the type of prober to use to monitor servers defined in this data center. The default value is inside-datacenter.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, glob, gtm link, gtm prober-pool, gtm server, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-04-24 gtm datacenter(1)

gtm distributed-app

NAME

distributed-app - Configures a Global Traffic Manager(tm) distributed application.

MODULE

gtm

SYNTAX

Configure the distributed-app component within the gtm module using the syntax in the following sections.

CREATE/MODIFY

```
create distributed-app [name]
modify distributed-app [name]
options:
  app-service [[string] | none]
  dependency-level [datacenter | link | none | server | wideip]
  description [string]
  disabled-contexts
    [add | delete | modify | replace-all-with] {
      [datacenter | link | server] [name] ...
    }
  disabled-contexts none
  persistence [enabled | disabled]
  persist-cidr-ipv4 [integer]
```

```
persist-cidr-ipv6 [integer]
ttl-persistence [integer]
wideips
  [add | delete | replace-all-with] {
    [name] ...
  }
wideips [default | none]
```

```
edit distributed-app
[ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

```
reset-stats distributed-app
reset-stats distributed-app
[ [ [name] | [glob] | [regex] ] ... ]
```

```
DISPLAY
list distributed-app
list distributed-app [ [ [name] | [glob] | [regex] ] ... ]
show running-config distributed-app
show running-config distributed-app
[ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  partition
```

```
show distributed-app
show distributed-app [name]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  detail
  field-fmt
```

```
DELETE
delete distributed-app [name]
```

DESCRIPTION

You can use the distributed-app component to create, modify, display, or delete a distributed application.

EXAMPLES

```
create distributed-app DA1
```

Creates a distributed application named DA1 with options set to the default values.

```
list distributed-app DA1 all-properties
```

Displays all properties of the distributed application named DA1.

OPTIONS

app-service

Specifies the name of the application service to which the distributed application belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the distributed application. Only the application service can modify or delete the distributed application.

dependency-level

Specifies the resources that must be in the available state before this distributed application is considered available. The options are:

datacenter

All of the data centers on the member list of this distributed application must be in an available state before the system considers the distributed application available.

link All of the links on the member list of this distributed application must be in an available state before the system considers the distributed application available.

none The distributed application has no dependencies. This value effectively disables this option. This is the default value.

server

All of the servers on the member list of this distributed application must be in an available state before the system considers the distributed application available.

wideip

All of the Wide IPs on the member list of this distributed application must be in an available state before the system considers the distributed application available.

description

User defined description.

disabled-contexts

Specifies the components that you want to add to or delete from this distributed application as disabled-contexts. You can also replace all of the components that are currently listed as disabled-contexts for

this distributed application with other components. The default value is none.

The possible values are:

datacenter

Specifies the datacenters, by name, to which the system does not send traffic from this distributed application.

link Specifies the links, by name, to which the system does not send traffic from this distributed application.

none There are no components to which the system does not send traffic from this distributed application. This value effectively disables this option.

server

Specifies the servers, by name, to which the system does not send traffic from this distributed application.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this object resides.

persistence

When enabled, if a local DNS server makes repetitive requests on behalf of a client, the system reconnects the client to the same resource as previous requests. The default value is disabled.

persist-cidr-ipv4

Specifies a mask used to group IPv4 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

persist-cidr-ipv6

Specifies a mask used to group IPv6 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ttl-persistence

Specifies, in seconds, the length of time for which the persistence entry is valid. The default value is 3600.

wideips

Specifies the Wide IPs, by name, that you want to add to or delete from this distributed application. You can also replace all of the Wide IPs that are currently associated with this distributed application with other Wide IPs. The default value is none.

A Wide IP is a collection of one or more domain names that maps to one or more groups of virtual servers managed either by BIG-IP(r) systems, or by host servers. The Global Traffic Manager load balances name resolution requests across the virtual servers that are defined in the Wide IP that is associated with the requested domain name.

SEE ALSO

create, delete, glob, gtm link, gtm server, create, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012. All rights reserved.

BIG-IP 2014-10-22 gtm distributed-app(1)

gtm global-settings general

NAME

general - Configures the general settings for the Global Traffic Manager.

MODULE

gtm global-settings

SYNTAX

Modify or display the general component within the gtm global-settings module using the syntax in the

following sections.

MODIFY

modify general

options:

- automatic-configuration-save-timeout [integer]
- auto-discovery [no | yes]
- auto-discovery-interval [integer]
- cache-ldns-servers [no | yes]
- domain-name-check [allow-underscore | none]
- drain-persistent-requests [no | yes]
- forward-status [enable | disable]
- gtm-sets-recursion [no | yes]
- heartbeat-interval [integer]
- iquery-cipher-list [string]
- iquery-crl-validation-depth [full | device]
- iquery-minimum-tls-version [string]
- iquery-reverify-on-crl-becoming-active [no | yes]
- iquery-reverify-on-crl-expiring [no | yes]
- iquery-reverify-on-crl-file-update [no | yes]
- iquery-use-expired-crls [no | yes]
- iquery-use-not-yet-active-crls [no | yes]
- iquery-use-revoked-certs [never | existing | always]
- monitor-disabled-objects [no | yes]
- nethsm-timeout [integer]
- nsec3-types-bitmap-strict [enable | disable]
- peer-leader [name]
- send-wildcard-rrs [enable | disable]
- static-persist-cidr-ipv4 [integer]
- static-persist-cidr-ipv6 [integer]
- synchronization [no | yes]
- synchronization-group-name [name]
- synchronization-time-tolerance [integer]
- synchronization-timeout [integer]
- synchronize-zone-files [no | yes]
- synchronize-zone-files-timeout [integer]
- topology-allow-zero-scores [no | yes]
- virtuals-depend-on-server-state [no | yes]
- wideip-zone-nameserver [string]

edit general

options:

- all-properties
- non-default-properties
- one-line

DISPLAY

list

list general

show running-config general

show running-config general [option name]

options:

- all-properties
- non-default-properties

DESCRIPTION

You can use the general component to modify or display the General Traffic Manager settings.

EXAMPLES

```
modify general auto-discovery no
```

Turns off auto-discovery for the Global Traffic Manager.

```
list general all-properties
```

Displays all properties of the general settings for the Global Traffic Manager.

OPTIONS

automatic-configuration-save-timeout

Sets the timeout, in seconds, indicating how long to wait after a GTM configuration change before automatically saving the GTM configuration to the `bigip_gtm.conf`. A timeout of `-1` will cause the GTM configuration to NEVER be saved. A value of `0` will cause the GTM configuration to be saved immediately. The default value is 15 seconds.

auto-discovery

Specifies whether the auto-discovery process is activated for this system. The default value is no.

auto-discovery-interval

Specifies the frequency, in seconds, between system attempts to discover network components. The default value is 30.

cache-ldns-servers

Specifies whether the system retains, in cache, all local DNS servers that make requests. The default value is yes.

You must enable this option if you want the system to store and use the LDNS path information.

domain-name-check

Specifies the parameters for the Global Traffic Manager to use when performing domain name checking. The default value is allow-underscore.

The possible values are:

allow-underscore

The Global Traffic Manager checks domain names according to the specifications in RFC 1123 Requirements for Internet Hosts - Application and Support, except that underscores are allowed.

none

No validation is performed. Anything is allowed.

idn-compatible

Deprecated since v12.1.0. Equivalent to allow-underscore. Value of idn-compatible will be saved as allow-underscore.

strict

Deprecated since v12.1.0. Equivalent to allow-underscore. Value of strict will be saved as allow-underscore.

drain-persistent-requests

Specifies, when set to yes, that when you disable a pool, load-balanced, persistent connections remain connected until the TTL expires. The default value is yes. If you set this option to no, any persistent connections terminate immediately when a pool is disabled.

forward-status

Specifies, when set to enabled, that the availability status change for GTM objects will be shared with subscribers. This option will enable iControl clients to receive event notifications when a change occurs.

gtm-sets-recursion

Specifies, when set to yes, that the system enables recursive DNS queries, regardless of whether the requesting local DNS enabled recursive queries. The default value is no.

heartbeat-interval

Specifies the frequency at which the Global Traffic Manager queries other BIG-IP(r) systems for updated data. When configuring monitors for BIG-IP systems, F5 Networks recommends that the probe-interval option for the monitor be equal to or greater than the this option. The default value is 10.

iquery-cipher-list

This is a ":" separated list of cipher specifications as accepted by the "openssl ciphers" command. OpenSSL will use the cipher list to negotiate a mutually acceptable cipher with the server during iQuery connection setup.

iquery-crl-validation-depth

Determines which CRL(s) are required during certificate validation for iQuery connections. The default value is full.

The possible values are:

full

A CRL must exist for every certificate authority in the certificate chain.

device

A CRL must exist for the certificate authority that issued the certificate. CRL(s) for other certificate authorities in the certificate chain are not used.

iquery-minimum-tls-version

This is a string to specify the minimum TLS version that will be offered by the client (GTM) during iQuery connection negotiation.

iquery-reverify-on-crl-becoming-active

Specifies, when set to yes, that all existing iQuery connections will have their certificates reverified whenever a whenever a CRL becomes active (thisUpdate is reached). The default value is yes.

iquery-reverify-on-crl-expiring

Specifies, when set to yes, that all existing iQuery connections will have their certificates reverified whenever a CRL expires (nextUpdate is reached). The default value is yes.

iquery-reverify-on-crl-file-update

Specifies, when set to yes, that all existing iQuery connections will have their certificates reverified whenever the CRL file is updated. The default value is yes.

iquery-send-wildcard-rrs

Specifies, when set to enable, that WildIPs or WildIP aliases that contain wildcards will autogenerate Resource Records in the BIND database. The default value is disable.

iquery-use-expired-crls

Specifies, when set to yes, that the validation of an iQuery SSL certificate can use an expired CRL (the \"nextUpdate\" field of the CRL in the past). The default value is yes.

iquery-use-not-yet-active-crls

Specifies, when set to yes, that the validation of an iQuery SSL certificate can use a not yet active CRL (the \"thisUpdate\" field of the CRL in the future). The default value is yes.

iquery-use-revoked-certs Specifies the action to take when a certificate is found to be revoked during the

verification of an iQuery connection.
The options are:

never

Do not allow the usage of revoked certificates. All new connections that are found to be revoked will be rejected. Any existing connections that are found to now be revoked will be disconnected.

existing

Only allow the usage of revoked certificates on previously established iQuery connections. Reject all new connections with certificates that are found to be revoked.

always

Allow the usage of revoked certificates on all new and existing iQuery connections.

monitor-disabled-objects

Specifies, when set to yes, that the system will continue to monitor objects even if the objects are disabled. The default value is no.

nethsm-timeout

Time to wait on a NetHSM key creation operation for DNSSEC before retry. Default is 20 seconds.

nsec3-types-bitmap-strict

When the nsec3-types-bitmap-strict setting has a default value of disabled the BIG-IP responds permissively to DS record queries when authenticating denial of existence. That is to say, the NSEC3 types bitmap will contain NS, even if we cannot be sure such a record exists.

When the setting is set to non-default value enabled (ie strict), the BIG-IP will only confirm the existence of the NS record (via the types bitmap of the NSEC3) when the zone is configured as an unsecured delegation on the DNSSEC Zone. If it is not configured, the BIG-IP will respond with TXT in the types bitmap.

peer-leader

Specifies the name of a GTM server to be used for executing certain features, such as creating DNSSEC keys.

send-wildcard-rrs

Specifies, when set to enable, that WildIPs or WildIP aliases that contain wildcards will autogenerate Resource Records in the BIND database. The default value is disable.

static-persist-cidr-ipv4

Specifies the number of bits of the IPv4 address that the system considers when using the Static Persist load balancing mode. The default value is 32.

static-persist-cidr-ipv6

Specifies the number of bits of the IPv6 address that the system considers when using the Static Persist load balancing mode. The default value is 128.

synchronization

Specifies whether this system is a member of a synchronization group. The default value is no.

Members of the synchronization group continuously share configuration and metrics collection information. The synchronization group can contain Global Traffic Managers and Link Controllers.

synchronization-group-name

Specifies the name of the synchronization group to which the system belongs. The default name is default.

synchronization-time-tolerance

Specifies the number of seconds that one system clock can be out of sync with another system clock, in the synchronization group. If the variance between the clock times is higher than the time tolerance setting, the system logs the time difference once per hour.

Possible values are 0 (zero), and 5 - 600. (Values 1 through 4 are automatically set to 5, and 0 (zero) turns time synchronization off.) The default value is 10 seconds.

Note: If you are using NTP to synchronize the clock with a time server, select a time tolerance other than 0 (zero). When you do this, the system uses the synchronization-time-tolerance option as a fail-over mechanism if NTP is disabled for any reason.

synchronization-timeout

Specifies the number of seconds that the system attempts to synchronize the Global Traffic Manager configuration with a synchronization group member. If the synchronization times out, the system tries again. The default value is 180.

synchronize-zone-files

Specifies whether the system synchronizes zone files among the synchronization group members. The default value is no.

synchronize-zone-files-timeout

Specifies the number of seconds that a synchronization group member attempts to synchronize its zone files with a synchronization group member. If the synchronization times out, the system tries again. The default value is 300.

topology-allow-zero-scores

Specifies if topology load-balancing or QoS load-balancing with topology enabled will return pool members with zero topology scores. The default value is yes.

virtuals-depend-on-server-state

Specifies whether the system marks a virtual server down when the server on which the virtual server is configured can no longer be reached via iQuery. The default value is yes.

wideip-zone-nameserver

Specifies the DNS Nameserver to use for all NS records for automatically generated DNS Zones created for all Wide IPs. It should be set to a registered DNS Nameserver for the Wide IPs.

SEE ALSO

edit, gtm global-settings load-balancing, gtm global-settings metrics, gtm global-settings metrics-exclusions, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2020-03-04 gtm global-settings general(1)

gtm global-settings load-balancing

NAME

load-balancing - Configures the load balancing settings for the Global Traffic Manager(tm).

MODULE

gtm global-settings

SYNTAX

Modify or display the load-balancing component within the gtm global-settings module using the syntax in the following sections.

MODIFY

modify load-balancing

options:

failure-rcode [noerror | formerr | servfail | nxdomain | notimpl | refused]
failure-rcode-ttl [integer]
failure-rcode-response [disabled | enabled]
ignore-path-ttl [no | yes]
respect-fallback-dependency [no | yes]
topology-longest-match [no | yes]
topology-prefer-edns0-client-subnet [disabled | enabled]
verify-vs-availability [no | yes]

edit load-balancing

options:

all-properties
non-default-properties

DISPLAY

list

list load-balancing

show running-config load-balancing

show running-config load-balancing [option]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the load-balancing component to modify or display the load balancing settings for the Global Traffic Manager.

EXAMPLES

modify load-balancing ignore-path-ttl yes

Specifies that dynamic load balancing methods can use path data, even after the time-to-live (TTL) for the path data expires.

list load-balancing all-properties

Displays all properties of the load balancing settings for the Global Traffic Manager.

OPTIONS

failure-rcode

Specifies the DNS RCODE used when failure-rcode-response is enabled. Default is noerror. Options include noerror (no type exists at this name), formerr (format error in query), servfail (unable to process query), nxdomain (name does not exist), notimpl (no support for this kind of query), and refused (refuse to process based on policy). If failure-rcode-ttl is non-zero, only the Authority section of the noerror

or nxdomain response will include a SOA record.

failure-rcode-response

When enabled, specifies that the system returns a RCODE response to Wide IP requests after exhausting all load-balancing methods. This affects all Wide IPs and may only be overridden by a more specific enabled configuration of a Wide IP. This response is an authoritative empty answer from the system for record requests. With this option enabled, the system responds faster to requests for which it does not have viable answers configured. The default value is disabled.

failure-rcode-ttl

Specifies the negative caching TTL of the SOA for the RCODE response. The default is 0, meaning no SOA is included (i.e. no caching).

ignore-path-ttl

Specifies, when set to yes, that dynamic load balancing methods can use path data, even after the time-to-live (TTL) for the path data expires. The default value is no.

respect-fallback-dependency

Specifies, when set to yes, that the system accepts virtual server status when the load balancing mode changes to the mode specified by the fallback-mode option of the pool. The default value is no.

topology-longest-match

Specifies, when set to yes, that the system evaluates all topology records in the topology statement, and then selects the topology record that most specifically matches the IP address in an LDNS request (in other words, has the longest match). When this option is set to no, the system selects the first record in the topology statement that matches the request.

topology-prefer-edns0-client-subnet

Specifies, when set to enabled, that the system should use the edns0 client subnet option (if one exists) instead of the source address when using topology load balancing. When this option is set to disabled or if the query did not contain a client subnet option, the system will fall back to the source address.

When disabled this option can be overridden by a per wide IP setting, gtm wideip [wideip type] [wideip name] topology-prefer-edns0-client-subnet [disabled | enabled].

verify-vs-availability

Specifies, when set to yes, that the system checks the availability of virtual servers before sending a connection to those virtual servers. The default value is no.

SEE ALSO

edit, gtm global-settings general, gtm global-settings metrics, gtm global-settings metrics-exclusions, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2016. All rights reserved.

BIG-IP 2017-09-28 gtm global-settings load-balancing(1)

gtm global-settings metrics-exclusions

NAME

metrics-exclusions - Configures the IP addresses that you want to exclude from the Global Traffic Manager(tm) metrics.

MODULE

gtm global-settings

SYNTAX

Modify or display the metrics-exclusions within the gtm global-settings module using the syntax in the following sections.

MODIFY

modify metrics-exclusions

options:

```
addresses [add | delete | none | replace-all-with] {  
  [ip address]...  
}
```

edit metrics-exclusions

options:

all-properties

DISPLAY

list

list metrics-exclusions

```
show running-config metrics-exclusions
  addresses
options:
  all-properties
  one-line
```

DESCRIPTION

You can use the metrics-exclusions component to exclude IP addresses from the Global Traffic Manager metrics.

EXAMPLES

```
modify metrics-exclusions addresses add {10.10.10.1}
```

Excludes the IP address 10.10.10.1 from inclusion in the Global Traffic Manager metrics.

```
list metrics-exclusions
```

Displays the IP addresses that are excluded from the Global Traffic Manager metrics.

OPTIONS

```
ip address
```

Specifies the IP addresses that you want to add to or delete from the exclusion list, or with which you want to replace all existing IP addresses that are currently on the exclusion list.

SEE ALSO

edit, gtm global-settings general, gtm global-settings load-balancing, gtm global-settings metrics, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 gtm global-settings metrics-exclusions(1)

gtm global-settings metrics

NAME

metrics - Configures metrics for the Global Traffic Manager(tm).

MODULE

gtm global-settings

SYNTAX

Modify or display the metrics component within the gtm global-settings module using the syntax in the following sections.

MODIFY

```
modify metrics
```

```
options:
```

```
  default-probe-limit [integer]
```

```
  hops-ttl [integer]
```

```
  hops-packet-length [integer]
```

```
  hops-sample-count [integer]
```

```
  hops-timeout [integer]
```

```
  inactive-lDNS-ttl [integer]
```

```
  lDNS-update-interval [integer]
```

```
  inactive-paths-ttl [integer]
```

```
  max-synchronous-monitor-requests [integer]
```

```
  metrics-caching [integer]
```

```
  metrics-collection-protocols none
```

```
  metrics-collection-protocols
```

```
    [add | delete | replace-all-with] {
```

```
  [dns-dot | dns-rev | icmp | tcp | udp] ...
```

```
  }
```

```
  path-ttl [integer]
```

```
  paths-retry [integer]
```

```
edit metrics
```

```
options:
```

```
  all-properties
```

```
  non-default-properties
```

```
  one-line
```

DISPLAY

```
list
```

```
list metrics
```

```
show running-config metrics
```

show running-config metrics [option]

options:

all-properties

non-default-properties

DESCRIPTION

You can use the metrics component to modify or display the Global Traffic Manager metrics settings.

EXAMPLES

modify metrics default-probe-limit 10

Sets the default probe limit for the Global Traffic Manager to 10.

list metrics all-properties

Displays all properties of the metrics settings for the Global Traffic Manager.

OPTIONS

default-probe-limit

Specifies the number of probe attempts that the system performs before removing the path from the metrics. The default value is 12.

hops-ttl

Specifies the number of seconds that the system considers traceroute utility data to be valid for name resolution and load balancing. The default value is 604800. Note that this option must be greater than the hops-timeout option.

hops-packet-length

Specifies the length of packets, in bytes, that the system sends to a local DNS server to determine the path information between the two systems. Valid values are 64 - 500. The default value is 64.

hops-sample-count

Specifies the number of packets that the system sends to a local DNS server to determine the path information between those two systems. Valid values are 1 - 10. The default value is 3.

hops-timeout

Specifies the number of seconds that the big3d daemon waits for a probe. Valid values are 1 - 10. The default value is 3.

inactive-ldns-ttl

Specifies the number of seconds that an inactive LDNS remains in the cache. Each time an LDNS makes a request, the clock starts again. Valid values are 60 - 31536000 (1 year). The default value is 2419200 (28 days).

ldns-update-interval

Specifies the number of seconds that a tmm will wait before sending an update for a LDNS which has been accessed. The default value is 20 seconds.

inactive-paths-ttl

Specifies the number of seconds that a path remains in the cache after its last access. Valid values are 60 - 31536000 (1 year). The default value is 604800 (7 days).

max-synchronous-monitor-requests

Specifies how many monitor requests are executed simultaneously. This value should only be changed if requested by F5 Support. The default value is 20.

metrics-caching

Specifies the interval (in seconds) at which the system dumps path and other metrics data. Valid values are 0 through 604800. The default value is 3600; 0 (zero) turns this feature off.

metrics-collection-protocols

Specifies the protocols that the system uses to collect metrics information relevant to LDNS servers.

path-ttl

Specifies the number of seconds that the system considers path data to be valid for name resolution and load balancing purposes. The default value is 2400. Note that this option must be greater than the paths-retry option and less than or equal to 2419200 (28 days).

paths-retry

Specifies the interval (in seconds) at which the system retries the path data. Valid values are 1 - 600 (10 minutes). The default value is 120.

SEE ALSO

edit, gtm global-settings general, gtm global-settings load-balancing, gtm global-settings metrics-exclusions, list, modify, show, tmmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013. All rights reserved.

gtm iquery

NAME

iquery - Displays information about iQuery.

MODULE

gtm

SYNTAX

Display the iquery component within the gtm module using the syntax in the following sections.

DISPLAY

show iquery

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DESCRIPTION

You can use the iquery component to display iQuery statistics.

EXAMPLES

show iquery

Displays iQuery statistics in the system default units.

show iquery field-fmt

Displays iQuery statistics in field format.

OPTIONS

For information about options for the command show, see show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2014-06-11 gtm iquery(1)

gtm ldns

NAME

ldns - Displays local domain name system (LDNS) statistics for the Global Traffic Manager(tm).

MODULE

gtm

SYNTAX

Display the ldns component within the gtm module using the syntax in the following section.

DISPLAY

show ldns

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DESCRIPTION

You can use the ldns component to display LDNS statistics.

EXAMPLES

show ldns

Displays LDNS statistics in the system default units.

show ldns field-fmt

Displays LDNS statistics in field format.

SEE ALSO
show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2014-06-11 gtm Idns(1)

gtm link

NAME
link - Configures Global Traffic Manager(tm) links.

MODULE
gtm

SYNTAX
Configure the link component within the gtm module using the syntax in the following sections.

CREATE/MODIFY
create link [name]
modify link [name]
options:
app-service [[string] | none]
cost-segments {
 { [up-to-bps [integer]] [dollars-per-mbps [integer]] }...
}
datacenter [string]
description [string]
[disabled | enabled]
duplex-billing [disabled | enabled]
limit-max-inbound-bps [integer]
limit-max-inbound-bps-status [disabled | enabled]
limit-max-outbound-bps [integer]
limit-max-outbound-bps-status [disabled | enabled]
limit-max-total-bps [integer]
limit-max-total-bps-status [disabled | enabled]
link-ratio [integer]
monitor [none | [name] [and [name]]...]
monitor min [integer] of { [name]... }
prepaid-segment [integer]
router-addresses
 [add | delete | modify | replace-all-with] {
[ip address] {
 app-service [[string] | none]
 description [string]
 device-name [[string] | none]
 translation [disabled | enabled]
 }
 }
 service-provider [name]
 uplink-address [ip address]
 weighting [price | ratio]

edit link [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line

reset-stats link
reset-stats link [[[name] | [glob] | [regex]] ...]

DISPLAY
list link
list link [[[name] | [glob] | [regex]] ...]
show running-config link
show running-config link [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

show link
show link [name]
options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

DELETE
delete link [name]

DESCRIPTION

You can use the link component to create, display, modify, or delete a link.

A link is a physical device that connects the network to the rest of the Internet. You can logically attach links to a collection of servers in order to manage access to the data sources on the network.

EXAMPLES

```
create link my_link datacenter DC1 router-addresses add {10.10.1.1}
```

Creates a link named my_link in the DC1 data center and adds the IP address of the router that uses this link.

```
list link non-default-properties
```

Displays all non-default properties for all links.

```
delete link my_link
```

Deletes the link named my_link.

```
show link my_link detail
```

Show the servers and virtual servers associated with my_link.

```
show link all detail
```

Show the servers and virtual servers associated with each respective link in the system.

OPTIONS

app-service

Specifies the name of the application service to which the link belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the link. Only the application service can modify or delete the link.

cost-segments

Specifies the cost of each incremental segment of bandwidth. This option is valid only when the weighting option is set to price. Note that you cannot modify the list, only replace all of the options in the list.

By default, the list is empty. The options are:

dollars-per-mps

Specifies the cost in dollars per megabytes per second. By default this value is not specified.

up-to-bps

Specifies the cost in dollars per bytes per second. By default this value is not specified.

datacenter

Specifies the data center to which the link belongs.

description

User defined description.

detail

The detail option is used with the show display command. This shows the server IP addresses and virtual servers associated with this link. A server can have multiple server IP addresses, however, only the server IP addresses that use this link will be displayed. Assignment information for servers and virtuals will not be displayed if there are not any servers or virtuals that use this link. By default, links are automatically matched to server IP addresses and virtual servers according to their IP addresses. Explicit links may also be defined. How this link was assigned is displayed under the Link Assignment column: auto means that the system automatically assigned this link, and explicit means that the link was explicitly set by the user.

[disabled | enabled]

Specifies whether the link and its resources are available for load balancing. The default value is enabled.

duplex-billing

Enables or disables duplex billing for this link. The default value is enabled. This option is valid only when the weighting option is set to price.

disabled

The internet service provider (ISP) that supplies this link bills for bandwidth usage based on the total amount of inbound plus outbound traffic on the link.

enabled

The ISP that supplies this link bills for bandwidth usage based on the maximum amount of either inbound or outbound traffic on the link (whichever is higher), rather than billing for bandwidth usage based on the total amount of inbound plus outbound traffic on the link.

glob Displays the items that match the glob expression. See help glob for a description of glob expression

syntax.

`limit-max-inbound-bps`

Specifies the threshold for inbound traffic on the link. The default value is 0 (zero).

`limit-max-inbound-bps-status`

Enables or disables the `limit-max-inbound-bps` option for this link. The default value is disabled.

`limit-max-outbound-bps`

Specifies the threshold for inbound traffic on the link. The default value is 0 (zero).

`limit-max-outbound-bps-status`

Enables or disables the `limit-max-outbound-bps` option for this link. The default value is disabled.

`limit-max-total-bps`

Specifies the threshold as a sum of inbound and outbound traffic on the link. The default value is 0 (zero).

`limit-max-total-bps-status`

Enables or disables the `limit-max-total-bps` option for this link. The default value is disabled.

`link-ratio`

Specifies the frequency at which the system sends traffic through the link. The default value is 1.

Important: When you set this option, you must also set the weighting option to `ratio`.

`monitor`

Specifies the health monitors that the system uses to determine whether this link is available for load balancing. Multiple monitors may be specified with the `and` keyword. The `min` keyword is used to specify the minimum number of monitors that must succeed for this link to be declared up. The default value is `none`.

`name` Specifies a unique name for the component. This option is required for the commands `create` and `modify`.

`prepaid-segment`

Specifies the amount of bandwidth for which the system is prepaid. This option is valid only when the weighting option is set to `price`. The default value is 0 (zero).

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`router-addresses`

Specifies the IP addresses of the routers that use this link. A router address can be associated with only one link. You can use the following options:

`app-service`

Specifies the name of the application service to which the link belongs. The default value is `none`.
Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the link. Only the application service can modify or delete the link.

`description`

User defined description.

`device-name`

Deprecated in v13.0.0 Specifies the name of the device associated with this link. Defaults to `none`.

`translation`

Specifies the address that the link uses for translation when communicating between the network and the Internet. The default value is `any6`.

`service-provider`

Specifies the Internet Service Provider (ISP) or vendor that supplies this link. This allows you to provide a name or description that helps identify this link out of your data center. This is an optional field.

`uplink-address`

Specifies the IP address the system uses to gather Simple Network Management Protocol (SNMP) metrics from the router interface. When you configure an uplink address, the system sends SNMP requests to the IP addresses configured using the `router-addresses` option for this link.

`weighting`

Specifies the weighting methodology the system uses to select a link to which to send traffic. The default value is `ratio`. The options are:

`price`

The system continuously checks the performance of each link, and sends traffic through the link with the best performance data.

`ratio`

The system uses the value that you set in the `link-ratio` option to determine the link to which to send traffic.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `gtm datacenter`, `gtm server`, `list`, `modify`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2016-11-18 gtm link(1)

gtm listener

NAME

listener - Configures a Global Traffic Manager(tm) listener.

MODULE

gtm

SYNTAX

Configure the listener component within the gtm module using the syntax in the following sections.

CREATE/MODIFY

```
create listener [name]
modify listener [name]
options:
  address [ip address]
  advertise [yes | no]
  app-service [[string] | none]
  auto-lasthop [default | enabled | disabled ]
  description [string]
  [disabled | enabled]
  fallback-persistence [none | [profile name] ]
  ip-protocol [tcp | udp]
  last-hop-pool [ [pool_name] | none]
  mask { [ipv4] | [ipv6] }
  persist [replace-all-with] {
[profile_name ... ] {
  default [no | yes]
}
}
  persist none
  pool [ [pool_name] | none]
  port [service port]
  profiles [add | delete | replace-all-with] {
[profile name ... ] {
  context [all | clientside | serverside]
}
}
  rules { [none | [rule_name ... ] }
  source-address-translation {
  options:
pool [ [pool_name] | none]
type [ automap | snat | none ]
}
  source-port [change | preserve | preserve-strict]
  translate-address [enabled | disabled]
  translate-port [enabled | disabled]
  vlans none
  vlans
  [ add | delete | replace-all-with ] {
[vlan name]...
}
  vlans-disabled
  vlans-enabled
```

```
edit listener [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

```
reset-stats listener
reset-stats listener [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list listener
list listener [name]
show running-config listener
show running-config listener [ [ [name] | [glob] | [regex] ] ... ]
options:
```

all-properties
non-default-properties
partition
show listener
show listener [name]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE
delete listener [name]

DESCRIPTION

You can use the listener component to create, display, modify, or delete a listener.

A listener is an object that listens for DNS queries. Listeners are defined for specific IP addresses, and are always associated with port 53.

Important: When you create, modify, or delete a listener, the system saves the running configuration in the stored configuration files.

EXAMPLES

```
create listener my_listener address 10.10.1.1 persist replace-all-with { source_addr }
```

Creates a listener named `my_listener` with an IP address of 10.10.1.1, which uses the source address persistence method.

```
modify listener my_listener profiles replace-all-with { dns }
```

Replaces the profiles associated with the listener `my_listener`.

Note: To replace the profile associated with a listener, you must enclose the name of the new profile in curly brackets.

```
list listener non-default-properties
```

Displays all non-default properties for all listeners.

```
delete listener my_listener
```

Deletes the listener named `my_listener`.

OPTIONS

address

Specifies the IP address on which the system listens. The system receives traffic sent to this IP address and processes it as needed. This option is required.

advertise

Specifies whether to advertise the listener address to surrounding routers. The options are `yes` or `no`. The default value is `no`.

app-service

Specifies the name of the application service to which the listener belongs. The default value is `none`.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the listener. Only the application service can modify or delete the listener.

context

Specifies that the protocol profile is either a clientside or serverside profile. If not specified, the default value is `all` for both sides.

description

User defined description.

(enabled | disabled)

Specifies the state of the listener. The default value is `enabled`.

Note: When you disable a listener, the listener no longer accepts new connection requests. However, it allows current connections to finish processing before going to a down state.

fallback-persistence

Specifies a fallback persistence profile for the listener to use when the default persistence profile is not available. The default value is `none`.

glob Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

ip-protocol

Specifies the protocol on which this listener receives network traffic. The options are `udp` or `tcp`. The default value is `udp`.

last-hop-pool

Specifies the name of the last hop pool that you want the listener to use to direct reply traffic to the last hop router. The default value is `none`.

mask Specifies the netmask for a network listener only. This setting is required for a network listener.

The netmask clarifies whether the host bit is an actual zero or a wildcard representation. The default

value is 255.255.255.255 for IPv4 or ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff for IPv6.

`name` Specifies a unique name for the component. This option is required for the commands `create` and `modify`.

`partition`

Displays the administrative partition within which the listener resides.

`persist`

Specifies a list of profiles separated by spaces that the listener uses to manage connection persistence. The default value is none.

To enable persistence, typically you specify a single profile. However, you can specify multiple profiles in conjunction with `iRules(r)` that define a persistence strategy based on incoming traffic. In the case of multiple profiles, the default option specifies which profile you want the listener to use if an `iRule` does not specify a persistence method. When you specify multiple profiles, the default value of the default property is `no`. You can set the value of the default property to `yes` for only one of the profiles.

`pool` Specifies a default pool to which you want the listener to automatically direct traffic. The default value is none.

`port` Specifies the service port on which the listener listens for connections. When you create a listener, the default value is 53 if no port number is specified.

`profiles`

Specifies the DNS, statistics and protocol profiles to use for this listener. When a listener is created, if a DNS profile is not specified, the generic "dns" profile is added. If a protocol profile is not specified, then the generic "tcp" profile is added for TCP and the "udp_gtm_dns" profile is added for UDP. A listener always has DNS and protocol profiles once it is created. Only a statistics profile can be added to or deleted from a listener.

The `replace-all-with` command replaces the profiles with the specified ones. The unspecified DNS and protocol profiles are not changed. If statistics profiles are not specified, the `replace-all-with` command removes the existing statistics profile from the listener. When the protocol is modified, if profiles are not specified, a default protocol profile is used. DNS and statistics profiles will not change.

`rules`

Specifies a list of `iRules`, separated by spaces, that customize the listener to direct and manage traffic. The default value is none.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an `@` sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`source-address-translation`

Specifies the type of source address translation enabled for the listener as well as the pool that the source address translation will use.

`pool` Specifies the name of a SNAT pool used by the specified listener.

`type` Specifies the type of source address translation associated with the specified listener.

The options are:

`automap`

Specifies the use of self IP addresses for listener source address translation.

`none` Specifies no source address translation to be used by the listener.

`snat` Specifies the use of a SNAT pool of translation addresses for listener source address translation.

`source-port`

Specifies whether the system preserves the source port of the connection. The default value is `preserve`.

The options are:

`change`

Obfuscates internal network addresses.

`preserve`

Preserves the source port of the connection.

`preserve-strict`

Use this value only for UDP under very special circumstances, such as `nPath` or `transparent` (that is, no translation of any other L3/L4 field), where there is a 1:1 relationship between virtual IP addresses and node addresses, or when clustered multi-processing (CMP) is disabled.

`translate-address`

Enables or disables address translation for the listener. Disable address translation for a listener if you want to use the listener to load balance connections to any address. This option is useful when the system is load balancing devices that have the same IP address. The default value is `disabled`.

`translate-port`

Enables or disables port translation. Disable port translation for a listener, if you want to use the listener to load balance connections to any service. The default value is `disabled`.

vans

Specifies a list of VLANs on which traffic is either disabled or enabled, based on whether the `vans-disabled` or `vans-enabled` option is specified.

vans-disabled

Specifies that traffic is not accepted by this listener on the VLANs specified in the `vans` option. This option is mutually exclusive with the `vans-enabled` option.

vans-enabled

Specifies that traffic is accepted by this listener on only the VLANs specified in the `vans` option. This option is mutually exclusive with the `vans-disabled` option.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `net vlan`, `net vlan-group`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 gtm listener(1)

gtm monitor bigip-link

NAME

`bigip-link` - Configures a BIG-IP(r) Link monitor.

MODULE

`gtm monitor`

SYNTAX

Configure the `bigip-link` component within the `gtm monitor` module using the syntax in the following sections.

CREATE/MODIFY

`create bigip-link [name]`

`modify bigip-link [name]`

options:

`app-service` [[string] | none]

`defaults-from` [name]

`description` [string]

`destination` [ip address]

`ignore-down-response` [enabled | disabled]

`interval` [integer]

`timeout` [integer]

`edit bigip-link [[[name] | [glob] | [regex]] ...]`

options:

`all-properties`

`non-default-properties`

DISPLAY

`list bigip-link`

`list bigip-link [[[name] | [glob] | [regex]] ...]`

`show running-config bigip-link`

`show running-config bigip-link [[[name] | [glob] | [regex]] ...]`

options:

`all-properties`

`non-default-properties`

`one-line`

`partition`

DELETE

`delete bigip-link [name]`

Note: You cannot delete default monitors.

DESCRIPTION

You can use the `bigip-link` component to configure a custom monitor, or you can use the default BIG-IP Link monitor that the Global Traffic Manager provides. This type of monitor acquires data captured through monitors managed by a BIG-IP Link Controller.

EXAMPLES

```
create bigip-link my_bigip-link defaults-from bigip_link
```

Creates a monitor named `my_bigip-link` that inherits properties from the default BIG-IP Link monitor.

list bigip-link

Displays the properties of all of the BIG-IP Link monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is bigip_link.

description

User defined description.

destination

Specifies the IP address of the resource that is the destination of this monitor. The default value is *.

Possible values are:

* Specifies to perform a health check on the IP address of the node.

IP address

Specifies to perform a health check on the IP address that you specify, route the check through the IP address of the associated node, and mark the IP address of the associated node up or down accordingly.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 30 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

create, delete, edit, glob, gtm link, list, ltm node, modify, regex, show, tmsh,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2012-10-19 gtm monitor bigip-link(1)

gtm monitor bigip

NAME

bigip - Configures a BIG-IP(r) monitor.

MODULE

gtm monitor

SYNTAX

Configure the bigip component within the gtm monitor module using the syntax in the following sections.

```
CREATE/MODIFY
create bigip [name]
modify bigip [name]
options:
  aggregate-dynamic-ratios [average-members | average-nodes | none |
    sum-members | sum-nodes]
  app-service [[string] | none]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  ignore-down-response [enabled | disabled]
  interval [integer]
  timeout [integer]

edit bigip [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

DISPLAY
list bigip
list bigip [ [name] | [glob] | [regex] ] ... ]
show running-config bigip
show running-config bigip [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

DELETE
delete bigip [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the bigip component in the gtm monitor module to configure a custom monitor, or you can use the default BIG-IP(r) monitor that the Global Traffic Manager(tm) provides. The BIG-IP monitor is both a health and performance monitor. This type of monitor acquires data captured through monitors managed by a BIG-IP Local Traffic Manager(tm).

You can monitor only the following components with a BIG-IP monitor:

- Global Traffic Manager server
- Global Traffic Manager virtual server
- Local Traffic Manager server
- Local Traffic Manager virtual server

EXAMPLES

```
create bigip my_bigip defaults-from bigip
```

Creates a monitor named my_bigip that inherits properties from the default BIG-IP monitor.

```
list bigip
```

Displays the properties of all of the BIG-IP monitors.

OPTIONS

aggregate-dynamic-ratios

Specifies the monitor's response to a query. By default, the BIG-IP monitor uses the gtm_score value as the vs_score for a Local Traffic Manager virtual server.

You can use this option to override the default behavior using the following values:

average-members

Specifies that the monitor uses the average of the dynamic ratio values of the pool members associated with the pools that are associated with the virtual server as a response to a query.

average-nodes

Specifies that the monitor uses the average value of all of the nodes associated with the pool members that are associated with the pools that are associated with the virtual server as a response to a query.

none This is the default value.

sum-members

Specifies that the monitor uses the sum of the pool members as a response to a query.

sum-nodes

Specifies that the monitor uses the sum of the dynamic ratios of all of the nodes as a response to a query.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot

modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is bigip.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the address and port supplied by a virtual server.

*:port

Specifies to perform a health check on the virtual server with the IP address supplied by the virtual server and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the virtual server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 90 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

create, delete, edit, glob, gtm pool, gtm server, list, modify, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2017-08-15 gtm monitor bigip(1)

gtm monitor external

NAME

external - Configures an external monitor.

MODULE

gtm monitor

SYNTAX

Configure the external component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create external [name]

modify external [name]
options:
args [[arguments] | none]
defaults-from [name]
description [string]
destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
ignore-down-response [enabled | disabled]
interval [integer]
probe-timeout [integer]
run [none | [external monitor]]
timeout [integer]
user-defined [[name] [value] | [name] none]

edit external [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY

list external
list external [[[name] | [glob] | [regex]] ...]
show running-config external
show running-config external [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE

delete external [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the external component to configure a custom monitor, or you can use the default external monitor that the Global Traffic Manager provides. You can use this type of monitor to monitor services using your own programs.

EXAMPLES

```
create external my_external defaults-from external
```

Creates a monitor named my_external that inherits properties from the default external monitor.

```
list external
```

Displays the properties of all of the external monitors.

OPTIONS

args Specifies any command-line arguments that the external program requires. The default value is none.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is external.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`run` Specifies the external monitor file to be executed by the external monitor. The default value is none.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

user-defined

Specifies any user-defined command-line arguments and variables that the external program requires. Use the following syntax to specify a user defined parameter.

```
modify external my_external user-defined my_param_name my_param_value
```

Use the following syntax to remove a user defined parameter.

```
modify external my_external user-defined my_param_name none
```

SEE ALSO

`create`, `delete`, `edit`, `glob`, `gtm pool`, `list`, `modify`, `regex`, `show`, `tmsh`,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2017-08-15 gtm monitor external(1)

gtm monitor firepass

NAME

`firepass` - Configures a FirePass(r) monitor.

MODULE

`gtm monitor`

SYNTAX

Configure the firepass component within the `gtm monitor` module using the syntax in the following sections.

CREATE/MODIFY

```
create firepass [name]
modify firepass [name]
options:
  app-service [[string] | none]
  cipherlist [list]
  concurrency-limit [integer]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  ignore-down-response [enabled | disabled]
  interval [integer]
  max-load-average [floating point value]
  password [none | [password] ]
  probe-timeout [integer]
  timeout [integer]
  username [name]
```

```
edit firepass [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list firepass
list firepass [ [ [name] | [glob] | [regex] ] ... ]
show running-config firepass
show running-config firepass [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

DELETE
delete firepass [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the firepass component to configure a custom monitor, or you can use the default FirePass monitor that the BIG-IP(r) Global Traffic Manager(tm) provides. The FirePass monitor is both a health and performance monitor.

For more information about configuring monitors, refer to the Configuration Guide for BIG-IP(r) Global Traffic Management.

EXAMPLES

```
create firepass my_firepass defaults-from firepass_gtm
```

Creates a monitor named my_firepass that inherits properties from the default FirePass monitor.

```
list firepass
```

Displays the properties of all of the FirePass monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

cipherlist

Specifies the list of ciphers for this monitor. The default value is HIGH:!ADH.

concurrency-limit

Specifies the maximum percentage of licensed connections currently in use under which the monitor marks the FirePass system up. The default value is 95.

For example, a value of 95 percent means that the monitor marks the FirePass system up until 95 percent of licensed connections are in use. When the number of in-use licensed connections exceeds 95 percent, the monitor marks the FirePass system down.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is firepass_gtm.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *:*

Possible values are:

: Specifies to perform a health check on the address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the address you supply.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

max-load-average

Specifies the number that the monitor uses to mark the FirePass system up or down. The system compares value of this option against a one-minute average of the FirePass system load. When the FirePass system-load average falls within the specified value, the monitor marks the FirePass system up. When the average exceeds the setting, the monitor marks the system down.

The default value is 12.0.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

password
Specifies the password, if the monitored target requires authentication. The default value is none.

probe-timeout
Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 90 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username
Specifies the username, if the monitored target requires authentication. The default value is gtmuser.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2017-08-15 gtm monitor firepass(1)

gtm monitor ftp

NAME

ftp - Configures a File Transfer Protocol (FTP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the ftp component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create ftp [name]
modify ftp [name]
options:
  debug [no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  filename [ [filename] | none]
  ignore-down-response [enabled | disabled]
  interval [integer]
  mode [passive | port]
  password [none | [password] ]
  probe-timeout [integer]
  timeout [integer]
  username [name]
```

```
edit ftp [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list ftp
```

```
list ftp [ [ [name] | [glob] | [regex] ] ... ]
show running-config ftp
show running-config ftp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
DELETE
delete ftp [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the ftp component to configure a custom monitor, or you can use the default FTP monitor that the Global Traffic Manager provides. This type of monitor verifies the FTP service by attempting to download a specific file to the /var/tmp directory on the system. Once downloaded successfully, the file is not saved.

EXAMPLES

```
create ftp my_ftp defaults-from ftp
```

Creates a monitor named my_ftp that inherits properties from the default FTP monitor.

```
list ftp
```

Displays the properties of all of the FTP monitors.

OPTIONS

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is ftp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

filename

Specifies the full path and file name of the file that the system attempts to download. The health check is successful if the system can download the file. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

mode Specifies the data transfer process (DTP) mode. The default value is passive. The options are:

passive

Specifies that the monitor sends a data transfer request to the FTP server. When the FTP server receives the request, the FTP server then starts and establishes the data connection.

port Specifies that the monitor starts and establishes the data connection with the FTP server.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

password
Specifies the password, if the monitored target requires authentication. The default value is none.

partition
Displays the administrative partition within which the component resides.

probe-timeout
Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username
Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor ftp(1)

gtm monitor gateway-icmp

NAME

gateway-icmp - Configures a Gateway Internet Control Message Protocol (ICMP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the gateway-icmp component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create gateway-icmp [name]

modify gateway-icmp [name]

options:

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

ignore-down-response [enabled | disabled]

interval [integer]

probe-attempts [integer]

probe-interval [integer]

probe-timeout [integer]

timeout [integer]

transparent [enabled | disabled]

edit gateway-icmp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list gateway-icmp

list gateway-icmp [[[name] | [glob] | [regex]] ...]

show running-config gateway-icmp

show running-config gateway-icmp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties
one-line
partition

DELETE
delete gateway-icmp [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the gateway-icmp component to configure a custom monitor, or you can use the default Gateway ICMP monitor that the Global Traffic Manager provides. You can use a Gateway ICMP type of monitor for a virtual server, a server (that is, all of the virtuals on a specified server), a pool member, a pool (that is, all of the pool members of a specified pool), or a link.

EXAMPLES

```
create gateway-icmp my_imcp defaults-from gateway_icmp
```

Creates a monitor named my_imcp that inherits properties from the default Gateway ICMP monitor.

```
list gateway-icmp
```

Displays the properties of all of the Gateway ICMP monitors.

OPTIONS

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is gateway_icmp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *:*

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool member (the gateway) is marked up or down accordingly.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

probe-attempts

Specifies the number of times the BIG-IP(r) system attempts to probe the host server, after which the BIG-IP system considers the host server down or unavailable. The default value is 3 attempts.

probe-interval

Specifies the frequency at which the BIG-IP system probes the host server. The default value is 1 second.

probe-timeout

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default

value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

SEE ALSO

create, delete, edit, glob, gtm link, gtm pool, gtm server, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2017-08-15 gtm monitor gateway-icmp(1)

gtm monitor gtp

NAME

gtp - Configures a GPRS Tunneling Protocol (GTP) monitor. This monitor operates over UDP.

MODULE

gtm monitor

SYNTAX

Configure the gtp component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create gtp [name]

modify gtp [name]

options:

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

ignore-down-response [enabled | disabled]

interval [integer]

probe-attempts [integer]

probe-interval [integer]

probe-timeout [integer]

protocol-version [integer]

timeout [integer]

edit gtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list gtp

list gtp [[[name] | [glob] | [regex]] ...]

show running-config gtp

show running-config gtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete gtp [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the gtp component to configure a custom monitor, or you can use the default GTP monitor that the Global Traffic Manager provides. This type of monitor verifies the GPRS tunneling service by attempting to send GTP Echo Requests to a pool, pool member, or virtual server, and verifying the receipt of a well-formed Echo Response packet. This monitor supports GTP version 1 and version 2 over UDP.

For more information about configuring monitors, refer to the Configuration Guide for BIG-IP(r) Global Traffic Management.

EXAMPLES

create gtp my_gtp defaults-from gtp

Creates a monitor named my_gtp that inherits properties from the default GTP monitor.

list gtp

Displays the properties of all of the GTP monitors.

OPTIONS

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is gtp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

probe-attempts

Specifies the number of times the BIG-IP system attempts to probe the host server, after which the BIG-IP system considers the host server down or unavailable. The default value is 3.

probe-interval

Specifies the frequency at which the BIG-IP system probes the host server. The default value is 1.

probe-timeout

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

protocol-version

Specifies the GTP protocol version used to perform the exchange. GTP version 1 and GTP version 2 are supported. The default is version 1.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a non-conforming Echo Reply, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

create, delete, edit, glob, gtm pool, gtm server, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

gtm monitor http

NAME

http - Configures a Hypertext Transfer Protocol (HTTP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the http component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create http [name]

modify http [name]

options:

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

ignore-down-response [enabled | disabled]

interval [integer]

password [none | [password]]

probe-timeout [integer]

recv [none | [string]]

recv-status-code [none | [string]]

reverse [enabled | disabled]

send [none | [string]]

timeout [integer]

transparent [enabled | disabled]

username [[name] | none]

edit http [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list http

list http [[[name] | [glob] | [regex]] ...]

show running-config http

show running-config http [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete http [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the http component to configure a custom monitor, or you can use the default HTTP monitor that the Global Traffic Manager provides. This type of monitor verifies the HTTP service by attempting to receive specific content from a Web page.

EXAMPLES

```
create http my_http defaults-from http
```

Creates a monitor named my_http that inherits properties from the default HTTP monitor.

```
list http
```

Displays the properties of all of the HTTP monitors.

OPTIONS

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is http.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

`IP address:port (with the transparent option enabled)`

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool member (the gateway) is marked up or down accordingly.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`ignore-down-response`

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

`interval`

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`partition`

Displays the administrative partition within which the component resides.

`password`

Specifies the password if the monitored target requires authentication. The default value is none.

`probe-timeout`

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

`recv` Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

`recv-status-code`

Specifies the status codes that the monitor looks for in the returned resource. The default value is none.

The status code can be a number within the range 100-999. Multiple values can be specified in a space separated, quoted string like "200 300". User can also specify wildcard values like "2XX 4xx 300". If you do not specify a value for send, recv and recv-status-code options, the monitor performs a simple service check and connect only.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`reverse`

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use this mode only if you configure both the send and recv options.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

`send` Specifies the text string that the monitor sends to the target object.

The default setting is GET /, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html.

Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

`transparent`

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, gtm pool, gtm server, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2019-08-29 gtm monitor http(1)

gtm monitor https

NAME

https - Configures a Hypertext Transfer Protocol over Secure Socket Layer (HTTPS) monitor.

MODULE

gtm monitor

SYNTAX

Configure the https component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create https [name]

modify https [name]

options:

cert [[cert list] | none]

cipherlist [string]

compatibility [enabled | disabled]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

ignore-down-response [enabled | disabled]

interval [integer]

key [[key] | none]

password [none | [password]]

probe-timeout [integer]

recv [none | [string]]

recv-status-code [none | [string]]

reverse [enabled | disabled]

send [none | [string]]

timeout [integer]

transparent [enabled | disabled]

username [[name] | none]

edit https [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list https

list https [[[name] | [glob] | [regex]] ...]

show running-config https

show running-config https [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete https [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the http component to configure a custom monitor, or you can use the default HTTPS monitor that the Global Traffic Manager provides. This type of monitor verifies the HTTPS service by attempting to receive specific content from a Web page protected by Secure Socket Layer (SSL) security.

EXAMPLES

create https my_https defaults-from https

Creates a monitor named my_https that inherits properties from the default HTTPS monitor.

list https

Displays the properties of all of the HTTPS monitors.

OPTIONS

cert Specifies a fully-qualified path for a client certificate that the monitor sends to the target SSL server. The default value is none.

cipherlist

Specifies the list of ciphers for this monitor. The default list DEFAULT:+SHA:+3DES:+kEDH:!EXPORT is located in the file gtm_base_monitors.conf.

compatibility

Specifies, when enabled, that the SSL options setting (in OpenSSL) is set to ALL. The default value is enabled.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is https.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool member (the gateway) is marked up or down accordingly.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

key Specifies the RSA private key if the monitored target requires authentication. The key must be surrounded by quotation marks, for example, key \"client.key\". Note that if you specify a key, you must also specify a value for the cert option. The default value is none.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-status-code

Specifies the status codes that the monitor looks for in the returned resource. The default value is none.

The status code can be a number within the range 100-999. Multiple values can be specified in a space separated, quoted string like "200 300". User can also specify wildcard values like "2XX 4xx 300". If you do not specify a value for send, recv and recv-status-code options, the monitor performs a simple service check and connect only.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reverse

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use the this mode only if you configure both the send and recv options.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

send Specifies the text string that the monitor sends to the target object. The default value is GET /, which retrieves a default HTML file for a web site.

To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html. Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2019-08-29 gtm monitor https(1)

gtm monitor imap

NAME

imap - Configures an Internet Message Access Protocol (IMAP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the imap component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create imap [name]

modify imap [name]

options:

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

folder [[name] | none]

ignore-down-response [enabled | disabled]

interval [integer]

password [none | [password]]

probe-timeout [integer]

timeout [integer]

transparent [enabled | disabled]

username [[name] | none]

edit imap [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties

DISPLAY

list imap

list imap [[[name] | [glob] | [regex]] ...]

show running-config imap

show running-config imap [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line
- partition

DELETE

delete imap [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the imap component to configure a custom monitor, or you can use the default IMAP monitor that the Global Traffic Manager provides. This type of monitor verifies IMAP by attempting to open a specified mail folder on a server. This monitor is similar to the POP3 monitor.

EXAMPLES

```
create imap my_imap defaults-from imap
```

Creates a monitor named my_imap that inherits properties from the default IMAP monitor.

```
list imap
```

Displays the properties of all of the IMAP monitors.

OPTIONS

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is imap.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

folder

Specifies the name of the folder on the IMAP server that the monitor tries to open. The default value is INBOX.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and

modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor imap(1)

gtm monitor ldap

NAME

ldap - Configures a Lightweight Directory Access Protocol (LDAP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the ldap component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create ldap [name]

modify ldap [name]

options:

base [none | [string]]

chase-referrals [no | yes]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

filter [[LDAP key] | none]

ignore-down-response [enabled | disabled]

interval [integer]

mandatory-attributes [no | yes]

password [none | [password]]

probe-timeout [integer]

security [none | ssl | tls]

timeout [integer]

username [[name] | none]

edit ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

```
list ldap
list ldap [ [ [name] | [glob] | [regex] ] ... ]
show running-config ldap
show running-config ldap [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete ldap [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the ldap component to configure a custom monitor, or you can use the default LDAP monitor that the Global Traffic Manager provides. This type of monitor verifies the LDAP service by attempting to authenticate the specified user.

EXAMPLES

```
create ldap my_ldap defaults-from ldap
```

Creates a monitor named my_ldap that inherits properties from the default LDAP monitor.

```
list ldap
```

Displays the properties of all of the LDAP monitors.

OPTIONS

base Specifies the location in the LDAP tree from which the monitor starts the health check. A sample value is dc=bigip-test,dc=net. The default value is none.

chase-referrals
Specifies whether the monitor upon receipt of an LDAP referral entry chases that referral. The default value is yes.

The options are:

no Specifies that the system will treat a referral entry as a normal entry and refrain from querying the remote LDAP server(s) pointed to by the referral entry.

yes Specifies that the system upon receiving any referral entry from the monitored LDAP server query, the system will then query the corresponding LDAP server(s) pointed to by the LDAP query. If the query for the referral is unsuccessful the system will mark the monitored LDAP server down.

debug
Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from
Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is ldap.

description
User defined description.

destination
Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the address and port supplied by a pool member.

*:port
Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port
Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

filter
Specifies an LDAP key for which the monitor searches. A sample value is objectclass=*. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax

ignore-down-response
Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval
Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

mandatory-attributes
Specifies whether the target must include attributes in its response to be considered up. The default value is no. The options are:

no Specifies that the system performs only a one-level search (based on the value of the filter option), and does not require that the target returns any attributes.

yes Specifies that the system performs a sub-tree search, and if the target returns no attributes, the target is considered down.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

password
Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout
Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

security
Specifies the secure communications protocol that the monitor uses to communicate with the target. The default value is none.

The options are:

none Specifies that the system does not use a security protocol for communications with the target.

ssl Specifies that the system uses the SSL protocol for communications with the target.

tls Specifies that the system uses the TLS protocol for communications with the target.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username
Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014, 2016. All rights reserved.

BIG-IP 2017-08-15 gtm monitor ldap(1)

gtm monitor mssql

NAME

mssql - Configures a Microsoft(r) Windows(r) Structured Query Language (MSSQL) monitor.

MODULE
gtm monitor

SYNTAX
Configure the mssql component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY
create mssql [name]
modify mssql [name]
options:
count [0 | 1]
database [[name] | none]
debug [no | yes]
defaults-from [name]
description [string]
destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
ignore-down-response [enabled | disabled]
interval [integer]
password [none | [password]]
probe-timeout [integer]
recv [none | [string]]
recv-column [none | [string]]
recv-row [none | [string]]
send [none | [string]]
timeout [integer]
username [[name] | none]

edit mssql [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list mssql
list mssql [[[name] | [glob] | [regex]] ...]
show running-config mssql
show running-config mssql [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete mssql [name]

Note: You cannot delete default monitors.

DESCRIPTION
You can use the mssql component to configure a custom monitor, or you can use the default MSSQL monitor that the Global Traffic Manager provides. This type of monitor verifies Microsoft Windows SQL-based services.

EXAMPLES
create mssql my_mssql defaults-from mssql

Creates a monitor named my_mssql that inherits properties from the default MSSQL monitor.

list mssql

Displays the properties of all of the MSSQL monitors.

OPTIONS
count
Specifies the number of instances for which the system keeps a connection open. By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open. With this option you can assign multiple instances to the database while reducing the overhead that multiple open connections can cause.

A value of 0 (zero), the default, keeps the connection open for all instances. A value of 1 opens a new connection for each instance. Any other positive value keeps the connection open for that many instances; for example, a value of 5 keeps the connection open for five instances of this monitor.

database
Specifies the name of the database with which the monitor attempts to communicate. The default value is none.

debug
Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the

/var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `mssql`.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

`recv` Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database. If you do not specify a value for both the `send` and `recv` options, the monitor performs a simple service check and connect only.

recv-column

Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the `send` and `recv` options. The default value is none.

recv-row

Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the `send` and `recv` options. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`send` Specifies the SQL query that the monitor sends to the target database, for example, `SELECT count(*) FROM mytable`.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the `recv` option and ignores the option even if not null.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `gtm pool`, `list`, `modify`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor mssql(1)

gtm monitor mysql

NAME

mysql - Configures a MySQL(r) monitor.

MODULE

gtm monitor

SYNTAX

Configure the mysql component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create mysql [name]

modify mysql [name]

options:

count [0 | 1]

database [[name] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

ignore-down-response [enabled | disabled]

interval [integer]

password [none | [password]]

probe-timeout [integer]

recv [none | [string]]

recv-column [none | [string]]

recv-row [none | [string]]

send [none | [string]]

timeout [integer]

username [[name] | none]

edit mysql [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list mysql

list mysql [[[name] | [glob] | [regex]] ...]

show running-config mysql

show running-config mysql [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete mysql [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the mysql component to configure a custom monitor, or you can use the default MySQL monitor that the Global Traffic Manager provides. This type of monitor verifies Microsoft(r) Windows(r) SQL-based services.

EXAMPLES

```
create mysql my_mysql defaults-from mysql
```

Creates a monitor named my_mysql that inherits properties from the default MySQL monitor.

```
list mysql
```

Displays the properties of all of the MySQL monitors.

OPTIONS

count

Specifies the number of instances for which the system keeps a connection open. By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open.

With this option you can assign multiple instances to the database while reducing the overhead that multiple open connections can cause.

A value of 0 (zero), the default, keeps the connection open for all instances. A value of 1 opens a new connection for each instance. Any other positive value keeps the connection open for that many instances; for example, a value of 5 keeps the connection open for five instances of this monitor.

database

Specifies the name of the database with which the monitor attempts to communicate. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is mysql.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-column

Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

recv-row

Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help regex for

a description of regular expression syntax.

send Specifies the SQL query that the monitor sends to the target database, for example, `SELECT count(*) FROM mytable`.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the `recv` option and ignores the option even if not null.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username
Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `gtm pool`, `list`, `modify`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor mysql(1)

gtm monitor nntp

NAME

`nntp` - Configures a Network News Transfer Protocol (NNTP) monitor.

MODULE

`gtm monitor`

SYNTAX

Configure the `nntp` component within the `gtm monitor` module using the syntax in the following sections.

CREATE/MODIFY

```
create nntp [name]
modify nntp [name]
options:
  debug [no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[.port] ] ]
  ignore-down-response [enabled | disabled]
  interval [integer]
  newsgroup [ [name] | none]
  password [none | [password] ]
  probe-timeout [integer]
  timeout [integer]
  username [ [name] | none]
```

```
edit nntp [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list nntp
list nntp [ [name] | [glob] | [regex] ] ... ]
show running-config nntp
show running-config nntp [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete nntp [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the nntp component to configure a custom monitor, or you can use the default NNTP monitor that the Global Traffic Manager provides. This type of monitor verifies the Usenet News protocol service by attempting to retrieve a newsgroup identification string from the server.

EXAMPLES

```
create nntp my_nntp defaults-from nntp
```

Creates a monitor named my_nntp that inherits properties from the default NNTP monitor.

```
list nntp
```

Displays the properties of all of the NNTP monitors.

OPTIONS

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is nntp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

newsgroup

Specifies the name of the newsgroup that you are monitoring, for example alt.car.mercedes. The default value is none.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor nntp(1)

gtm monitor none

NAME

none - A base level NULL monitor used to indicate that no monitoring will be performed.

MODULE

gtm monitor

SYNTAX

The none component within the gtm monitor module can be used only with the syntax described in the following sections.

DISPLAY

```
list none
list none [ [ [name] | [glob] | [regex] ] ... ]
show running-config none
show running-config none [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

CREATE/MODIFY/DELETE

Note: You cannot create a custom monitor based off of the none monitor nor can you modify or delete the existing default none monitor.

DESCRIPTION

The none component is used to configure a GTM pool member that explicitly does not inherit its parent pool's monitor and remains un-monitored, regardless of the monitor setting on the parent pool.

EXAMPLES

```
create gtm pool a test_pool monitor http members add { test:vs_1 { monitor none } test:vs_2 { } }
```

Creates a gtm pool with an http monitor that has two members, vs_1 and vs_2. The pool member vs_1 is configured with the none monitor so it does not inherit http monitoring from its parent pool, whereas pool member vs_2 will inherit and use its parent's http monitor.

OPTIONS

Note: While the following options are specified as part of the definition of the none monitor, they cannot be changed since the none monitor is used to define the absence of a monitor, not an actual monitor.

app-service

Displays the application service to which the object belongs. The permanent value is none. Note: Regardless if the strict-updates option is enabled, you cannot modify or delete the object; these are unavailable options for the none monitor.

defaults-from

Specifies the name of the monitor from which this monitor inherits settings. The default value is none.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *:6666, also shown as *:ircu-2.

ignore-down-response

Specifies whether this monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 0 seconds.

partition

Displays the administrative partition within which the component resides. The default value is Common.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 0 seconds.

time-until-up

Specifies the amount of time in seconds after the first successful response before a node is marked up. A value of 0 causes a node to be marked up immediately after a valid response is received from the node. The default value is 86500.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 0 seconds.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0, which specifies that the system uses the value of the interval option whether the resource is up or down.

SEE ALSO

gtm pool, list, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015. All rights reserved.

BIG-IP 2017-08-29 gtm monitor none(1)

gtm monitor oracle

NAME

oracle - Configures an Oracle(r) monitor.

MODULE

gtm monitor

SYNTAX

Configure the oracle component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create oracle [name]
modify oracle [name]
options:
  count [0 | 1]
  database [ [name] | none]
  debug [no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  ignore-down-response [enabled | disabled]
  interval [integer]
  password [none | [password] ]
  probe-timeout [integer]
  recv [none | [string] ]
  recv-column [none | [string] ]
  recv-row [none | [string] ]
  send [none | [string] ]
  timeout [integer]
  username [ [name] | none]
```

```
edit oracle [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list oracle
list oracle [ [ [name] | [glob] | [regex] ] ... ]
show running-config oracle
show running-config oracle [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

one-line
partition

DELETE
delete oracle [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the oracle component to configure a custom monitor, or you can use the default Oracle monitor that the Global Traffic Manager provides. This type of monitor verifies services based on Oracle by attempting to perform an Oracle login to a service.

EXAMPLES

```
create oracle my_oracle defaults-from oracle
```

Creates a monitor named my_oracle that inherits properties from the default Oracle monitor.

```
list oracle
```

Displays the properties of all of the Oracle monitors.

OPTIONS

count

Specifies the number of instances for which the system keeps a connection open. By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open. With this option you can assign multiple instances to the database while reducing the overhead that multiple open connections can cause.

A value of 0 (zero), the default, keeps the connection open for all instances. A value of 1 opens a new connection for each instance. Any other positive value keeps the connection open for that many instances; for example, a value of 5 keeps the connection open for five instances of this monitor.

database

Specifies the name of the database with which the monitor attempts to communicate. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is oracle.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-column

Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

recv-row

Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

send Specifies the SQL query that the monitor sends to the target database, for example, SELECT count(*) FROM mytable.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor oracle(1)

gtm monitor pop3

NAME

pop3 - Configures a Post Office Protocol version 3 (POP3) monitor.

MODULE

gtm monitor

SYNTAX

Configure the pop3 component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create pop3 [name]

modify pop3 [name]

options:

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

ignore-down-response [enabled | disabled]

interval [integer]

password [none | [password]]

probe-timeout [integer]
timeout [integer]
username [[name] | none]

edit pop3 [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list pop3
list pop3 [[[name] | [glob] | [regex]] ...]
show running-config pop3
show running-config pop3 [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete pop3 [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the pop3 component to configure a custom monitor, or you can use the default POP3 monitor that the Global Traffic Manager provides. This type of monitor verifies the POP3 service by attempting to connect to a pool, pool member, or virtual server, log on as the specified user, and log off.

EXAMPLES

```
create pop3 my_pop3 defaults-from pop3
```

Creates a monitor named my_pop3 that inherits properties from the default POP3 monitor.

```
list pop3
```

Displays the properties of all of the POP3 monitors.

OPTIONS

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is pop3.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and

modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor pop3(1)

gtm monitor postgresql

NAME

postgresql - Configures a PostgreSQL(r) monitor.

MODULE

gtm monitor

SYNTAX

Configure the postgresql component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create postgresql [name]

modify postgresql [name]

options:

count [0 | 1]

database [[name] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

ignore-down-response [enabled | disabled]

interval [integer]

password [none | [password]]

probe-timeout [integer]

recv [none | [string]]

recv-column [none | [string]]

recv-row [none | [string]]

send [none | [string]]

timeout [integer]

username [[name] | none]

edit postgresql [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list postgresql

```
list postgresql [ [ [name] | [glob] | [regex] ] ... ]
show running-config postgresql
show running-config postgresql [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

DELETE
delete postgresql [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the postgresql component to configure a custom monitor, or you can use the default PostgreSQL monitor that the Global Traffic Manager provides. This type of monitor verifies Microsoft(r) Windows(r) SQL-based services.

EXAMPLES

```
create postgresql my_postgresql defaults-from postgresql
```

Creates a monitor named my_postgresql that inherits properties from the default PostgreSQL monitor.

```
list postgresql
```

Displays the properties of all of the PostgreSQL monitors.

OPTIONS

count

Specifies the number of instances for which the system keeps a connection open. By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open. With this option you can assign multiple instances to the database while reducing the overhead that multiple open connections can cause.

A value of 0 (zero), the default, keeps the connection open for all instances. A value of 1 opens a new connection for each instance. Any other positive value keeps the connection open for that many instances; for example, a value of 5 keeps the connection open for five instances of this monitor.

database

Specifies the name of the database with which the monitor attempts to communicate. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is postgresql.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

password
Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout
Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-column
Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

recv-row
Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

send Specifies the SQL query that the monitor sends to the target database, for example, SELECT count(*) FROM mytable.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username
Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor postgresql(1)

gtm monitor radius-accounting

NAME
radius-accounting - Configures a RADIUS accounting monitor for the BIG-IP(r) Global Traffic Manager.

MODULE
gtm monitor

SYNTAX
Configure the radius-accounting component within the gtm monitor module using the syntax shown in the following sections.

CREATE/MODIFY
create radius-accounting [name]
modify radius [name]
options:

app-service [[string] | none]
check-until-up [disabled | enabled]
debug [no | yes]
defaults-from [[name] | none]
description [string]
destination [ip address]
interval [integer]
manual-resume [disabled | enabled]
nas-ip-address [ip address]
secret [string]
time-until-up [integer]
timeout [integer]
username [none | [string]]

edit radius-accounting [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list radius-accounting

list radius-accounting [[[name] | [glob] | [regex]] ...]

show running-config radius

show running-config radius [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

DELETE

delete radius-accounting [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the radius-accounting component to configure a custom monitor, or you can use the default RADIUS accounting monitor that the Global Traffic Manager provides. This type of monitor provides information about the usage of the RADIUS service for accounting purposes.

EXAMPLES

```
create radius-accounting my_radius_acct defaults-from radius_accounting
```

Creates a monitor named my_radius_acct that inherits properties from the default RADIUS accounting monitor.

```
list radius-accounting
```

Displays the properties of all of the RADIUS accounting monitors.

OPTIONS

app-service

Specifies the name of the application service to which this monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this monitor. Only the application service can modify or delete this monitor.

check-until-up

When enabled, specifies that when an active and passive (inband) monitor are combined in an AND type of rule, the active monitor performs health checks only when the pool member is down, or until the pool member is marked as up. When the passive monitor marks the pool member down, the active monitor resumes health checks.

The default value is disabled.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is radius.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

***:port**

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

nas-ip-address

Specifies the network access server IP address that the system uses to identify itself to the RADIUS server. Using this option, multiple BIG-IP(r) systems can appear as a single network access device to the RADIUS server. The default value is none.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Specifies the secret the monitor needs to communicate with the resource. The default value is none.

time-until-up

Specifies the amount of time in seconds after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2014. All rights reserved.

BIG-IP 2016-06-29 gtm monitor radius-accounting(1)

gtm monitor radius

NAME

radius - Configures a Remote Access Dial-in User Service (RADIUS) monitor.

MODULE

gtm monitor

SYNTAX

Configure the radius component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create radius [name]
modify radius [name]
options:
  app-service [[string] | none]
  debug [no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  ignore-down-response [enabled | disabled]
  interval [integer]
  nas-ip-address [ [ip address] | none]
  password [none | [password] ]
  probe-timeout [integer]
  secret [none | [secret] ]
  timeout [integer]
  username [ [name] | none]
```

```
edit radius [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list radius
list radius [ [ [name] | [glob] | [regex] ] ... ]
show running-config radius
show running-config radius [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete radius [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the radius component to configure a custom monitor, or you can use the default RADIUS monitor that the Global Traffic Manager provides. This type of monitor verifies the RADIUS service by attempting to authenticate the specified user.

EXAMPLES

```
create radius my_radius defaults-from radius
```

Creates a monitor named my_radius that inherits properties from the default RADIUS monitor.

```
list radius
```

Displays the properties of all of the RADIUS monitors.

OPTIONS

app-service

Specifies the name of the application service to which this monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this monitor. Only the application service can modify or delete this monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is radius.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

***:port**

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

nas-ip-address

Specifies the network access server IP address that the system uses to identify itself to the RADIUS server. Using this option, multiple BIG-IP(r) systems can appear as a single network access device to the RADIUS server. The default value is none.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Specifies the secret the monitor needs to communicate with the resource. The default value is none.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor radius(1)

gtm monitor real-server

NAME

real-server - Configures a RealServer(r) monitor.

MODULE

gtm monitor

SYNTAX

Configure the real-server component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create real-server [name]
```

```
modify real-server [name]
```

options:

```
defaults-from [name]
```

```
description [string]
```

```
ignore-down-response [enabled | disabled]
```

```
interval [integer]
```

```
metrics [ [metrics] | none]
```

```
probe-timeout [integer]
```

```
timeout [integer]
```

```
edit real-server [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list real-server
```

```
list real-server [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config real-server
```

```
show running-config real-server [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
agent
```

```
all-properties
```

```
command
```

```
method
```

```
non-default-properties
```

```
one-line
```

```
partition
```

DELETE

```
delete real-server [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the real-server component to configure a custom monitor, or you can use the default RealServer monitor that the Global Traffic Manager provides. This type of monitor checks the performance of a pool, pool member, or virtual server that is running the RealServer data collection agent, and then dynamically load balances traffic accordingly.

EXAMPLES

```
create real-server my_real-server defaults-from real_server
```

Creates a monitor named my_real-server that inherits properties from the default RealServer monitor.

```
list real-server
```

Displays the properties of all of the RealServer monitors.

OPTIONS

agent

Displays the agent for the monitor. The default agent is Mozilla/4.0 (compatible: MSIE 5.0; Windows NT). You cannot modify the agent.

command

Displays the command that the system uses to obtain the metrics from the resource. See the documentation for this resource for information on available commands. You cannot modify the command.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is real_server.

description

User defined description.

glob

Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

method

Displays the GET method. You cannot modify the method.

metrics

Specifies the performance metrics that the commands collect from the target. The default value is ServerBandwidth:1.5, CPUPercentUsage, MemoryUsage, TotalClientCount.

partition
Displays the administrative partition within which the component resides.

probe-timeout
Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO
create, delete, edit, glob, gtm pool, gtm server, list, ltm node, modify, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2010-07-02 gtm monitor real-server(1)

gtm monitor scripted

NAME
scripted - Configures a Scripted monitor.

MODULE
gtm monitor

SYNTAX
Configure the scripted component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY
create scripted [name]
modify scripted [name]
options:
debug [no | yes]
defaults-from [name]
description [string]
destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
filename [[filename] | none]
ignore-down-response [enabled | disabled]
interval [integer]
probe-timeout [integer]
timeout [integer]

edit scripted [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list scripted
list scripted [[[name] | [glob] | [regex]] ...]
show running-config scripted
show running-config scripted [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete scripted [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the scripted component to configure a custom monitor, or you can use the default scripted monitor that the Global Traffic Manager provides.

EXAMPLES

```
create scripted my_scripted defaults-from scripted
```

Creates a monitor named my_scripted that inherits properties from the default Scripted monitor.

```
list scripted
```

Displays the properties of all of the scripted monitors.

OPTIONS

`debug`

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

`no` Specifies that the system does not redirect error messages and additional information related to this monitor.

`yes` Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is scripted.

`description`

User defined description.

`destination`

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

`filename`

Specifies the name of a file in the `/config/eav/` directory on the system. The user-created file contains the send and expect data that the monitor uses for the monitor check. The default value is none.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`ignore-down-response`

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

`interval`

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`partition`

Displays the administrative partition within which the component resides.

`probe-timeout`

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the

system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor scripted(1)

gtm monitor sip

NAME

sip - Configures a Session Initiation Protocol (SIP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the sip component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create sip [name]
modify sip [name]
options:
  cert [ [cert list] | none]
  cipherlist [list]
  compatibility [enabled | disabled]
  debug [ no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  filter [any | none | status]
  filter-neg [any | none | status]
  headers [ [new line separated headers] | none]
  ignore-down-response [enabled | disabled]
  interval [integer]
  key [ [key] | none]
  mode [sips | tcp | tls | udp]
  probe-timeout [integer]
  request [none | [string] ]
  username [ [name] | none]
```

```
edit sip [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list sip
list sip [ [ [name] | [glob] | [regex] ] ... ]
show running-config sip
show running-config sip [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete sip [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the sip component to configure a custom monitor, or you can use the default SIP monitor that the Global Traffic Manager provides. This type of monitor checks the status of SIP Call-ID services on a device. The SIP protocol enables real-time messaging, voice, data, and video.

EXAMPLES

```
create sip my_sip defaults-from sip
```

Creates a monitor named my_sip that inherits properties from the default SIP monitor.

```
list sip
```

Displays the properties of all of the SIP monitors.

OPTIONS

`cert` Specifies a fully-qualified path for a client certificate that the monitor sends to the target SSL server. The default value is none.

`cipherlist`
Specifies the list of ciphers for this monitor. The default value is none.

`compatibility`
Specifies, when enabled, that the SSL options setting (in OpenSSL) is set to ALL. The default value is enabled.

`debug`
Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

`no` Specifies that the system does not redirect error messages and additional information related to this monitor.

`yes` Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

`defaults-from`
Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is sip.

`description`
User defined description.

`destination`
Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`
Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`
Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

`filter`
Specifies the SIP status codes that the target can return to be considered up. By default the system always accepts status codes whose value is in the 100s, 200s, or 300s.

The options are:

`any` Specifies that the monitor accepts any SIP status codes.

`none` Specifies that the monitor does not accept any other SIP status codes. This is the default value.

`status`
Specifies one or more status codes that you want to add to the monitor.

`filter-neg`
Specifies the SIP status codes that the target can return to be considered down. By default the system always accepts status codes according to `sip-monitor.filter`. After checking that, the status code is checked against this key. If a code is also in `sip-monitor.filter`, the node is marked up.

The options are:

`any` Specifies that the monitor rejects all SIP status codes that are not in `sip-monitor.filter`.

`none` Specifies that the monitor does not specifically reject any other SIP status codes. This is the default value.

`status`
Specifies one or more status codes that you want to add to the monitor.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`headers`
Specifies the set of SIP headers in the SIP message that is sent to the target. Separate each header with a new line. The default value is none.

`ignore-down-response`
Specifies whether the monitor ignores a down response from the system it is monitoring. The default value

is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

key Specifies the key if the monitored target requires authentication. The default value is none.

mode Specifies the transport protocol that the monitor uses to communicate with the target. The default mode is udp. The options are:

sips Specifies that the monitor uses SIPS to communicate with the target.

tcp Specifies that the monitor uses TCP to communicate with the target.

tls Specifies that the monitor uses TLS to communicate with the target, and the SIP URI is SIPS.

udp Specifies that the monitor uses UDP to communicate with the target.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

request

Specifies the SIP request line in the SIP message that is sent to the target. The default value is none.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014, 2016. All rights reserved.

BIG-IP 2019-10-14 gtm monitor sip(1)

gtm monitor smtp

NAME

smtp - Configures a Simple Mail Transport Protocol (SMTP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the smtp component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create smtp [name]
modify smtp [name]
options:
  debug [no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  domain [ [name] | none]
  ignore-down-response [enabled | disabled]
  interval [integer]
  probe-timeout [integer]
```

timeout [integer]

edit smtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list smtp

list smtp [[[name] | [glob] | [regex]] ...]

show running-config smtp

show running-config smtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete smtp [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the smtp component to configure a custom monitor, or you can use the default SMTP monitor that the Global Traffic Manager provides. This type of monitor checks the status of a pool, pool member, or virtual server by issuing standard SMTP commands.

EXAMPLES

```
create smtp my_smtp defaults-from smtp
```

Creates a monitor named my_smtp that inherits properties from the default SMTP monitor.

```
list smtp
```

Displays the properties of all of the SMTP monitors.

OPTIONS

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is smtp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

domain

Specifies the domain name to check, for example, bigipinternal.com. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

probe-timeout
Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO
create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014, 2016. All rights reserved.

BIG-IP 2017-08-15 gtm monitor smtp(1)

gtm monitor snmp-link

NAME
snmp-link - Configures a Simple Network Management Protocol (SNMP) link monitor.

MODULE
gtm monitor

SYNTAX
Configure the snmp-link component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY
create snmp-link [name]
modify snmp-link [name]
options:
app-service [[string] | none]
community [[name] | none]
defaults-from [name]
description [string]
destination [ip address]
ignore-down-response [enabled | disabled]
interval [integer]
port [[integer] | none]
probe attempts [integer]
probe-interval [integer]
probe-timeout [integer]
timeout [integer]
version [[integer] | none]

edit snmp-link [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list snmp-link
list snmp-link [[[name] | [glob] | [regex]] ...]
show running-config snmp-link
show running-config snmp-link [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete snmp-link [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the snmp-link component to configure a custom monitor, or you can use the default SNMP Link monitor that the Global Traffic Manager provides. This type of monitor checks the current CPU, memory, and disk usage of a pool, pool member, or virtual server that is running an SNMP data collection agent, and then dynamically load balances traffic accordingly.

EXAMPLES

```
create snmp-link my_snmp-link defaults-from snmp_link
```

Creates a monitor named my_snmp-link that inherits properties from the default SNMP Link monitor.

```
list snmp-link
```

Displays the properties of all of the SNMP Link monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

community

Specifies the community name that the BIG-IP(r) system must use to authenticate with the host server through SNMP. The default value is public.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is snmp_link.

description

User defined description.

destination

Specifies the IP address of the resource that is the destination of this monitor. The default value is *.

Possible values are:

* Specifies to perform a health check on the IP address of the node.

IP address

Specifies to perform a health check on the IP address that you specify, route the check through the IP address of the associated node, and mark the IP address of the associated node up or down accordingly.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

port

Specifies the port number to which this monitor sends SNMP traps. The default value is 161.

probe-attempts

Specifies the number of times the BIG-IP system attempts to probe the host server, after which the BIG-IP system considers the host server down or unavailable. The default value is 3.

probe-interval

Specifies the frequency at which the BIG-IP system probes the host server. The default value is 0.

probe-timeout

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 30 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

version
Specifies the SNMP version the monitor uses. The default value is none.

SEE ALSO
create, delete, edit, glob, list, ltm node, modify, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2012-10-19 gtm monitor snmp-link(1)

gtm monitor snmp

NAME
snmp - Configures a Simple Network Management Protocol (SNMP) monitor.

MODULE
gtm monitor

SYNTAX
Configure the snmp component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY
create snmp [name]
modify snmp [name]
options:
app-service [[string] | none]
community [[name] | none]
defaults-from [name]
description [string]
destination [[ipv4 address[:port]] | [ipv6 address[.port]]]
ignore-down-response [enabled | disabled]
interval [integer]
port [integer]
probe attempts [integer]
probe-interval [integer]
probe-timeout [integer]
timeout [integer]
version [[integer] | none]

edit snmp [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list snmp
list snmp [[[name] | [glob] | [regex]] ...]
show running-config snmp
show running-config snmp [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete snmp [name]

Note: You cannot delete default monitors.

DESCRIPTION
You can use the snmp component to configure a custom monitor, or you can use the default SNMP monitor that the Global Traffic Manager provides. The SNMP monitor is both a health and performance monitor. This type of monitor checks the performance of a server running an SNMP agent such as UC Davis, for the purpose of load balancing traffic to that server.

EXAMPLES
create snmp my_snmp defaults-from snmp_gtm

Creates a monitor named `my_snmp` that inherits properties from the default SNMP monitor.

`list snmp`

Displays the properties of all of the SNMP monitors.

OPTIONS

`app-service`

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

`community`

Specifies the community name that the BIG-IP(r) system must use to authenticate with the host server through SNMP. The default value is `public`.

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `snmp_gtm`.

`description`

User defined description.

`destination`

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`ignore-down-response`

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

`interval`

Specifies the frequency at which the system issues the monitor check. The default value is 90 seconds.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`port` Specifies the port number to which this monitor sends SNMP traps. The default value is 161.

`probe-attempts`

Specifies the number of times the BIG-IP system attempts to probe the host server, after which the BIG-IP system considers the host server down or unavailable. The default value is 1.

`probe-interval`

Specifies the frequency at which the BIG-IP system probes the host server. The default value is 1.

`probe-timeout`

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 180 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

`version`

Specifies the SNMP version the monitor uses. The default value is `v1`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `gtm pool`, `list`, `modify`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2017-08-15 gtm monitor snmp(1)

gtm monitor soap

NAME

soap - Configures a Simple Object Access Protocol (SOAP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the soap component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create soap [name]

modify soap [name]

options:

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

expect-fault [no | yes]

ignore-down-response [enabled | disabled]

interval [integer]

method [string]

namespace [[name] | none]

parameter-name [[name] | none]

parameter-type [bool | int | long | [string]]

parameter-value [none | [integer] | [string]]

password [none | [password]]

probe-timeout [integer]

protocol [[none] | [protocol]]

return-type [bool | char | double | int | long | short | [string]]

return-value [none | [integer] | [string]]

soap-action [string]

timeout [integer]

url-path [none | [string]]

username [[name] | none]

edit soap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list soap

list soap [[[name] | [glob] | [regex]] ...]

show running-config soap

show running-config soap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete soap [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the soap component to configure a custom monitor, or you can use the default SOAP monitor that the Global Traffic Manager provides. This type of monitor tests a Web service based on SOAP.

EXAMPLES

```
create soap my_soap defaults-from soap
```

Creates a monitor named my_soap that inherits values from the system default SOAP monitor.

```
list soap
```

Displays the properties of all of the SOAP monitors.

OPTIONS

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the type of monitor you want to use to create the new monitor. The default value is soap.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

expect-fault

Specifies whether the value of the method option causes the monitor to expect a SOAP fault message. The default value is no.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

method

Specifies the method by which the monitor contacts the resource.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

namespace

Specifies the name space for the Web service you are monitoring, for example, `http://example.com/`. The default value is none.

parameter-name

If the method has a parameter, specifies the name of that parameter. The default value is bool.

parameter-type

Specifies the parameter type. The default value is none.

parameter-value

Specifies the value for the parameter. The default value is none.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

protocol

Specifies the protocol that the monitor uses to communicate with the target. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

return-type

Specifies the type for the returned parameter. The default value is bool.

return-value

Specifies the value for the returned parameter. The default value is none.

soap-action

Specifies the value for the SOAPAction header. The default value is the empty string.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

url-path

Specifies the URL for the Web service that you are monitoring, for example, /services/mysevice.aspx. The default value is none.

username

Specifies the user name if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor soap(1)

gtm monitor tcp-half-open

NAME

tcp-half-open - Configures a Transmission Control Protocol (TCP) Half Open monitor.

MODULE

gtm monitor

SYNTAX

Configure the tcp-half-open component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create tcp-half-open [name]

modify tcp-half-open [name]

options:

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

ignore-down-response [enabled | disabled]

interval [integer]

probe-attempts [integer]

probe-interval [integer]

probe-timeout [integer]

timeout [integer]

transparent [disabled | enabled]

edit tcp-half-open [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tcp-half-open

list tcp-half-open [[[name] | [glob] | [regex]] ...]

show running-config tcp-half-open

show running-config tcp-half-open [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete tcp-half-open [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the tcp-half-open component to configure a custom monitor, or you can use the default TCP Half Open monitor that the Global Traffic Manager provides.

For more information about configuring monitors, refer to the Configuration Guide for BIG-IP(r) Global Traffic Management.

EXAMPLES

```
create tcp-half-open my_tcp-half-open defaults-from tcp_half_open
```

Creates a monitor named my_tcp-half-open that inherits properties from the default TCP Half Open monitor.

```
list tcp-half-open
```

Displays the properties of all of the TCP Half Open monitors.

OPTIONS

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `tcp_half_open`.

`description`

User defined description.

`destination`

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

`IP address:port (with the transparent option enabled)`

Specifies to perform a health check on the server at the IP address and port you specify, route the check through the IP address and port supplied by the pool member, and mark the pool member (the gateway) up or down accordingly.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`ignore-down-response`

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

`interval`

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`probe-attempts`

Specifies the number of times the BIG-IP system attempts to probe the host server, after which the BIG-IP system considers the host server down or unavailable. The default value is 3.

`probe-interval`

Specifies the frequency at which the BIG-IP system probes the host server. The default value is 1.

`probe-timeout`

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

SEE ALSO

create, delete, edit, glob, gtm pool, gtm server, list, modify, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2017-08-15 gtm monitor tcp-half-open(1)

gtm monitor tcp

NAME

tcp - Configures a Transmission Control Protocol (TCP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the tcp component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create tcp [name]

modify tcp [name]

options:

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

ignore-down-response [enabled | disabled]

interval [integer]

probe-timeout [integer]

recv [none | [string]]

reverse [enabled | disabled]

send [none | [string]]

timeout [integer]

transparent [disabled | enabled]

edit tcp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tcp

list tcp [[[name] | [glob] | [regex]] ...]

show running-config tcp

show running-config tcp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete tcp [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the tcp component to configure a custom monitor, or you can use the default TCP monitor that the Global Traffic Manager provides.

EXAMPLES

```
create tcp my_tcp defaults-from tcp
```

Creates a monitor named my_tcp that inherits properties from the default TCP monitor.

```
list tcp
```

Displays the properties of all of the TCP monitors.

OPTIONS

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is tcp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port you specify, route the check through the IP address and port supplied by the pool member, and mark the pool member (the gateway) up or down accordingly.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reverse

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use this mode only if you configure both the send and recv options.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

send Specifies the text string that the monitor sends to the target object. The default setting is GET /, which retrieves a default HTML file for a web site.

To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html. Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool

members through firewalls. The default value is disabled.

SEE ALSO

create, delete, edit, glob, gtm pool, gtm server, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2017-08-15 gtm monitor tcp(1)

gtm monitor udp

NAME

udp - Configures a User Datagram Protocol (UDP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the udp component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create udp [name]

modify udp [name]

options:

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

ignore-down-response [enabled | disabled]

interval [integer]

probe-attempts [integer]

probe-interval [integer]

probe-timeout [integer]

recv [none | [string]]

reverse [enabled | disabled]

send [none | [string]]

timeout [integer]

transparent [disabled | enabled]

edit udp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list udp

list udp [[[name] | [glob] | [regex]] ...]

show running-config udp

show running-config udp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete udp [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the udp component to configure a custom monitor, or you can use the default UDP monitor that the Global Traffic Manager provides. This type of monitor verifies the UDP service by attempting to send UDP packets to a pool, pool member, or virtual server, and receiving a reply.

For more information about configuring monitors, refer to the Configuration Guide for BIG-IP(r) Global Traffic Management.

EXAMPLES

```
create udp my_udp defaults-from udp
```

Creates a monitor named my_udp that inherits properties from the default UDP monitor.

```
list udp
```

Displays the properties of all of the UDP monitors.

OPTIONS

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `udp`.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port you specify, route the check through the IP address and port supplied by the pool member, and mark the pool member (the gateway) up or down accordingly.

glob Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

name Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

partition

Displays the administrative partition within which the component resides.

probe-attempts

Specifies the maximum number of times the BIG-IP system will attempt to probe an unresponsive host server before marking the server as down/unavailable. The default value is 3 attempts. This attribute requires the usage of a Receive String, otherwise only 1 probe attempt occurs.

probe-interval

Specifies the time between individual probe attempts sent by the BIG-IP to the host server. The default value is 1 second. This attribute requires the usage of a Receive String, otherwise only 1 probe attempt occurs.

probe-timeout

Specifies the number of seconds after which the BIG-IP system times out the probe request to the BIG-IP system. The default value is 5 seconds.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names. If you do not specify a value for both the `send` and `recv` options, the monitor performs a simple service check and connect only.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

reverse

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

send Specifies the text string that the monitor sends to the target object. The default value is "default send string".

To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html. Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the value of the recv option and ignores the option even if not null.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent
Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

SEE ALSO

create, delete, edit, glob, gtm pool, gtm server, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014. All rights reserved.

BIG-IP 2018-02-13 gtm monitor udp(1)

gtm monitor wap

NAME

wap - Configures a Wireless Application Protocol (WAP) monitor.

MODULE

gtm monitor

SYNTAX

Configure the wap component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY

create wap [name]

modify wap [name]

options:

accounting-node [none | [RADIUS server name]]
accounting-port [[integer] | none]
call-id [none | [RADIUS server 11 digit phone number]]
check-until-up [enabled | disabled]
debug [no | yes]
defaults-from [name]
description [string]
destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
framed-address [none | [RADIUS framed IP address]]
ignore-down-response [enabled | disabled]
interval [integer]
probe-timeout [integer]
recv [none | [string]]
secret [none | [password]]
send [none | [string]]
server-id [none | [RADIUS NAS-ID]]
session-id [none | [RADIUS session ID]]
timeout [integer]

edit wap [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list wap

list wap [[[name] | [glob] | [regex]] ...]

```
show running-config wap
show running-config wap[ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
DELETE
delete wap [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the wap component to configure a custom monitor, or you can use the default WAP monitor that the Global Traffic Manager provides. This type of monitor requests the URL specified in the send option, and finds the string specified in the recv option somewhere in the data returned by the URL response.

EXAMPLES

```
create wap my_wap defaults-from wap
```

Creates a monitor named my_wap that inherits properties from the default WAP monitor.

```
list wap
```

Displays the properties of all of the WAP monitors.

OPTIONS

accounting-node

Specifies the RADIUS server that provides authentication for the WAP target. Note that if you configure the accounting-port option, but you do not configure the this option, the system assumes that the RADIUS server and the WAP server are the same system.

accounting-port

Specifies the port that the monitor uses for RADIUS accounting. The default value is none. A value of 0 (zero) disables RADIUS accounting.

call-id

Specifies the 11-digit phone number for the RADIUS server. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is wap.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

framed-address

Specifies the RADIUS framed IP address. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

probe-timeout
Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

recv Specifies the text string that the monitor looks for in the returned resource. The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names. If you do not specify both a value for both the send and recv options, the monitor performs a simple service check and connect only. The default value is none.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret
Specifies the password the monitor needs to communicate with the resource. The default value is none.

send Specifies the text string that the monitor sends to the target object. The default setting is GET /, which retrieves a default HTML file for a web site.

To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html. Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if it is not null. The default value is none.

server-id
Specifies the RADIUS NAS-ID for this system when configuring a RADIUS server. The default value is none.

session-id
Specifies the RADIUS session identification number when configuring a RADIUS server. The default value is none.

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014. All rights reserved.

BIG-IP 2017-08-15 gtm monitor wap(1)

gtm monitor wmi

NAME
wmi - Configures a Windows(r) Management Instrumentation (WMI) monitor.

MODULE
gtm monitor

SYNTAX
Configure the wmi component within the gtm monitor module using the syntax in the following sections.

CREATE/MODIFY
create wmi [name]
modify wmi [name]
options:
command [[command] | none]

defaults-from [name]
description [string]
ignore-down-response [enabled | disabled]
interval [integer]
metrics [[integer] | none]
password [none | [password]]
probe-timeout [integer]
timeout [integer]
url [none | [URL]]
username [[name] | none]

edit wmi [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list wmi

list wmi [[[name] | [glob] | [regex]] ...]

show running-config wmi

show running-config wmi [[[name] | [glob] | [regex]] ...]

options:

agent
all-properties
method
non-default-properties
one-line
partition
post

DELETE

delete wmi [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the wmi component to configure a custom monitor, or you can use the default WMI monitor that the Global Traffic Manager provides. This type of monitor checks the performance of a pool, pool member, or virtual server that is running the WMI data collection agent, and then dynamically load balances traffic accordingly.

EXAMPLES

```
create wmi my_wmi defaults-from wmi
```

Creates a monitor named my_wmi that inherits properties from the default WMI monitor.

```
list wmi
```

Displays the properties of all of the WMI monitors.

OPTIONS

agent

Displays the agent for the monitor. The default agent is Mozilla/4.0 (compatible: MSIE 5.0; Windows NT). You cannot modify the agent.

command

Specifies the command that the system uses to obtain the metrics from the resource. See the documentation for this resource for information on available commands.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is wmi.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-down-response

Specifies whether the monitor ignores a down response from the system it is monitoring. The default value is disabled.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 30 seconds.

method

Displays the GET method. You cannot modify the method.

metrics

Specifies the performance metrics that the commands collect from the target. The default value is LoadPercentage, DiskUsage, PhysicalMemoryUsage:1.5, VirtualMemoryUsage:2.0.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

post Specifies the mechanism that the monitor uses for posting. The default value is RespFormat=HTML. You cannot change the post format for WMI monitors.

probe-timeout

Specifies the number of seconds after which the BIG-IP(r) system times out the probe request to the BIG-IP system. The default value is 5 seconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 120 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

url Specifies the URL that the monitor uses. The default value is /scripts/f5Isapi.dll.

username

Specifies the user name if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, gtm pool, list, ltm node, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-02-21 gtm monitor wmi(1)

gtm path

NAME

path - Displays path statistics for the Global Traffic Manager(tm).

MODULE

gtm

SYNTAX

Show the path statistics using the syntax in the following section.

DISPLAY

show path

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DESCRIPTION

You can use the path component to display path statistics for the Global Traffic Manager.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2017-04-05 gtm path(1)

gtm persist

NAME

persist - Displays persistence records for the Global Traffic Manager(tm).

MODULE

gtm

SYNTAX

Display statistics for the persist component within the gtm module using the syntax in the following section.

DISPLAY

show persist

options:

destination [[name] | none]

key [ip address | string]

level [application | wideip]

max-results [integer]

target-name [[name] | none]

target-type [datacenter | link | pool-member | server]

DESCRIPTION

You can use the persist component to display various persistence records based on the filtering options that you use.

EXAMPLES

show persist

Displays all Global Traffic Manager persistence records.

show persist level wideip

Displays persistence records only for wideip persistence.

OPTIONS

destination

Displays persistence records for the specified destination.

key Displays persistence records for the specified LDNS address or generic key.

level

Displays persistence records for the specified level (destination type), either wideip or application.

max-results

Specifies the maximum number of persistence records that you want the system to return.

target-name

Displays persistence records for the specified target name.

If target-type pool-member is specified, then a valid target-name filter for a terminal pool member is the full path virtual server name, for example /Common/vs-dns-1. Do not include the server name in the target-name filter field.

For a non-terminal member the valid target-name filter is the non-pathed DNS name, for example www.alt.example.com.

Please refer to the EXAMPLES section below for more info.

target-type

Displays persistence records for the specified type of target.

EXAMPLES

All persistence records:

```
[root@bigip-1:Active:Standalone] config # tms show gtm persist
```

```
Value Level:Destination -> Target Expiration
```

```
10.2.4.4 wideip:A:/Common/www.example.com -> pool-member:www.alt.example.com 11-03 17:23:54
```

```
10.2.4.1 wideip:A:/Common/www.example.com -> pool-member:/Common/gtm-server-1:/Common/vs-dns-1 11-03 17:23:54
```

```
10.2.4.3 wideip:A:/Common/www.example.com -> pool-member:/Common/gtm-server-1:/Common/vs-dns-2 11-03 17:23:54
```

```
10.2.4.5 wideip:A:/Common/www.example.com -> pool-member:www.alt.example.com 11-03 17:23:54
```

```
10.2.4.2 wideip:A:/Common/www.example.com -> pool-member:www.alt.example.com 11-03 17:23:54
```

```
Total persistence records returned: 5
```

Filter: show only records for pool member /Common/vs-dns-2

```
[root@bigip-1:Active:Standalone] config # tms show gtm persist target-type pool-member target-name /Common/vs-dns-2
```

```
Value Level:Destination -> Target Expiration
```

```
10.2.4.2 wideip:A:/Common/www.example.com -> pool-member:/Common/gtm-server-1:/Common/vs-dns-2 11-03 16:01:03
```

```
Total persistence records returned: 1
```

Notice that only the virtual server name /Common/vs-dns-2 is specified in the target-name filter.

Filter: show only records for pool member CNAME www.alt.example.com

```
[root@bigip-1:Active:Standalone] config # tms show gtm persist target-type pool-member target-name www.alt.example.com
```

```
Value Level:Destination -> Target Expiration
```

```
10.2.4.4 wideip:A:/Common/www.example.com -> pool-member:www.alt.example.com 11-03 17:23:54
```

```
10.2.4.5 wideip:A:/Common/www.example.com -> pool-member:www.alt.example.com 11-03 17:23:54
```

```
10.2.4.2 wideip:A:/Common/www.example.com -> pool-member:www.alt.example.com 11-03 17:23:54
```

Total persistence records returned: 3

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013. All rights reserved.

BIG-IP 2015-11-06 gtm persist(1)

gtm pool a

NAME

a - Configures A load balancing pools for the Global Traffic Manager(tm).

MODULE

gtm pool

SYNTAX

Modify the Global Traffic Manager pool a component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

create pool a [name]

modify pool a [name]

options:

alternate-mode [drop-packet | fallback-ip | global-availability
| none | packet-rate | ratio | return-to-dns | round-robin
| static-persistence | topology | virtual-server-capacity
| virtual-server-score]
app-service [[string] | none]
description [string]
[disabled | enabled]
dynamic ratio [disabled | enabled]
fallback-ip [ip address]
fallback-mode [completion-rate | cpu | drop-packet | fallback-ip
| fewest-hops | global-availability | kilobytes-per-second
| least-connections | lowest-round-trip-time | none
| packet-rate | quality-of-service | ratio | return-to-dns
| round-robin | static-persistence | topology
| virtual-server-capacity | virtual-server-score]
limit-max-bps [integer]
limit-max-bps-status [disabled | enabled]
limit-max-connections [integer]
limit-max-connections-status [disabled | enabled]
limit-max-pps [integer]
limit-max-pps-status [disabled | enabled]
load-balancing-mode [completion-rate | cpu | drop-packet
| fallback-ip | fewest-hops | global-availability
| kilobytes-per-second | least-connections
| lowest-round-trip-time | packet-rate | quality-of-service
| ratio | return-to-dns | round-robin | static-persistence
| topology | virtual-server-capacity | virtual-server-score]
manual-resume [disabled | enabled]
max-answers-returned [integer]
members none
members
[add | delete | modify | replace-all-with] {
[server-name:vs-name] {
options:
app-service [[string] | none]
depends-on none
depends-on
[add | delete | modify | replace-all-with] {
[server-name:vs-name]...
}
description [string]
[disabled | enabled]
limit-max-bps [integer]
limit-max-bps-status [disabled | enabled]
limit-max-connections [integer]
limit-max-connections-status [disabled | enabled]
limit-max-pps [integer]
limit-max-pps-status [disabled | enabled]
member-order [integer]

```

monitor [none | [name] [and [name] ]... ]
monitor min [integer] of { [name]... }
monitor require [integer] from [integer] { [name] }
ratio [integer]
}...
}
metadata none
metadata
  [add | delete | modify | replace-all-with] {
    [metadata_name ... ] {
app-service [[string] | none]
persist [ true | false ]
value [ "value content" ]
    }
  }
}
monitor [none | [name] [and [name] ]... ]
monitor min [integer] of { [name]... }
monitor require [integer] from [integer] { [name] }
qos-hit-ratio [integer]
qos-hops [integer]
qos-kilobytes-second [integer]
qos-lcs [integer]
qos-packet-rate [integer]
qos-rtt [integer]
qos-topology [integer]
qos-vs-capacity [integer]
qos-vs-score [integer]
ttl [integer]
verify-member-availability [disabled | enabled]

edit pool a [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

reset-stats pool a
reset-stats pool a [ [ [name] | [glob] | [regex] ] ... ]

DISPLAY
list pool a
list pool a [ [ [name] | [glob] | [regex] ] ... ]
show running-config pool a
show running-config pool a [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  members vs-name
  one-line
  partition

show pool a
show pool a [name]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  detail
  field-fmt

DELETE
delete pool a [name]

```

Note: You must remove all references to a pool before you can delete the pool.

DESCRIPTION

You can use the pool component to configure the A pool definitions on the Global Traffic Manager. You use a pool to group member servers together to use a common load balancing algorithm.

EXAMPLES

```

create pool a mypool members add {
member myServer:myVs
}
monitor http

```

Creates a Global Traffic Manager A pool with one member myServer:myVs using the Round Robin load balancing method, and default HTTP monitor checks for member availability.

```
delete pool a my_pool
```

Deletes the pool my_pool.

```
show pool a
```

Displays statistics for all A pools.

```
list pool a my_pool
```

Displays settings of pool my_pool.

OPTIONS

alternate-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred method is unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option. The default value is round-robin.

The options are:

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fallback-ip

Specifies that the Global Traffic Manager returns the IP address that you specify as an answer to the query.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

none Specifies that the system skips the alternate load balancing mode and immediately tries the load balancing mode specified in the fallback-mode option.

Note that if the value of the fallback-mode option is none, and you have multiple pools configured, the Global Traffic Manager uses the next available pool.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

app-service

Specifies the name of the application service to which this pool belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool. Only the application service can modify or delete this pool.

description

User defined description.

[disabled | enabled]

Specifies whether this pool is available for load balancing. The default value is enabled.

dynamic-ratio

Enables or disables a dynamic ratio load balancing algorithm for this pool. This option is applicable only when you also configure the load-balancing-mode option for the pool with one of these dynamic load balancing modes: completion-rate, fewest-hops, kilobytes-per-second, least-connections, lowest-round-trip-times, quality-of-service, virtual-server-capacity, or virtual-server-score.

When this option is disabled (the default), the system uses only the server or virtual server with the best metrics, or highest quality of service (QoS) score, for load balancing. When dynamic-ratio is enabled, the system treats QoS scores as ratios, and it uses each server or virtual server in proportion to the ratio determined by the QoS calculation.

fallback-ip

Specifies the IPv6 address of the server to which the system directs requests in the event that the load balancing methods configured for this pool fail to return a valid virtual server. The default value is ::.

fallback-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the

members of this pool, if the preferred and alternate modes are unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option, and the alternate mode using the alternate-mode option. The default value is return-to-dns.

The options are:

`completion-rate`

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

`cpu` Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

`drop-packet`

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

`fallback-ip`

Specifies that the Global Traffic Manager returns the IP address that you specify as an answer to the query.

`fewest-hops`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

`global-availability`

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

`kilobytes-per-second`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

`least-connections`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

`lowest-round-trip-time`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

`none` Specifies that there is no fallback mode. If the system cannot use the preferred or alternate load balancing modes, it uses the next pool to resolve the request. If there are no more pools available, the result is the same as when the value for the fallback-mode option is return-to-dns.

`packet-rate`

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

`quality-of-service`

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

`ratio`

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

`return-to-dns`

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

`round-robin`

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

`static-persistence`

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

`topology`

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

`virtual-server-capacity`

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

`virtual-server-score`

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

limit-max-bps

Specifies the maximum allowable data throughput rate, in bits per second, for the virtual servers in the pool. If the network traffic volume exceeds this value, the system marks the pool as unavailable. The default value is 0 (zero).

limit-max-bps-status

Enables or disables the limit-max-bps option for this pool. The default value is disabled.

limit-max-connections

Specifies the number of current connections allowed for the virtual servers in the pool. If the current connections exceed this value, the system marks the pool as unavailable. The default value is 0 (zero).

limit-max-connections-status

Enables or disables the limit-max-connections option for this pool. The default value is disabled.

limit-max-pps

Specifies the maximum allowable data transfer rate, in packets per second, for the virtual servers in the pool. If the network traffic volume exceeds this value, the system marks the pool as unavailable. The default value is 0 (zero).

limit-max-pps-status

Enables or disables the limit-max-pps option for this pool. The default value is disabled.

load-balancing-mode

Specifies the preferred load balancing mode that the system uses to load balance name resolution requests among the members of this pool. The default value is round-robin.

The options are:

completion-rate

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

cpu Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fallback-ip

Specifies that the Global Traffic Manager returns the IP address that you specify as an answer to the query.

fewest-hops

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

manual-resume

Enables or disables the manual resume function for this pool. If you leave this option disabled (the default), then a member of this pool automatically becomes available for load balancing when its status changes from down to up. When the manual-resume option is enabled, if the status of a member of this pool changes from up to down, the pool member remains disabled indefinitely until you manually re-enable it.

max-answers-returned

Specifies the maximum number of available virtual servers that the system lists in a response. The default value is 1.

members

Specifies the vs-name of the pool members. The default value is none.

You can also use the following options with pool members:

app-service

Specifies the name of the application service to which this pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool member. Only the application service can modify or delete this pool member.

depends-on

Specifies the name of the virtual server on which this pool member depends.

description

User defined description.

[enabled | disabled]

Specifies whether this pool member is available for load balancing. The default value is enabled.

limit-max-bps

Specifies the maximum allowable data throughput rate, in bits per second, for the pool member. If the network traffic volume exceeds this value, the system marks the pool member as unavailable.

limit-max-bps-status

Enables or disables the limit-max-bps option for this pool member. The default value is disabled.

limit-max-connections

Specifies the number of current connections allowed for this pool member. If the current connections exceed this value, the system marks this pool member as unavailable.

limit-max-connections-status

Enables or disables the limit-max-connection option for this pool member. The default value is disabled.

limit-max-pps

Specifies the maximum allowable data transfer rate, in packets per second, for this pool member. If the network traffic volume exceeds this value, the system marks this pool member as unavailable.

limit-max-pps-status

Enables or disables the limit-max-pps option for this pool member. The default value is disabled.

member-order

Specifies the order number of the pool member. The system uses this number with load balancing methods that involve prioritizing pool members, such as the Ratio load balancing method.

monitor

Specifies the health monitors that the system uses to determine whether it can use this pool member for load balancing. Multiple monitors may be specified with the and keyword. The min keyword is used to specify the minimum number of monitors that must succeed for this pool member to be declared up. The require keyword is used to specify the minimum number of probes that must succeed for this server to be declared up and the number of probes that should be used. The default value is none.

ratio

Specifies the weight of the pool member for load balancing purposes.

vs-name

Displays the name of the corresponding virtual server.

metadata

Associates user defined data, each of which has name and value pair and persistence. Persistent(default) means the data will be saved into config file.

monitor

Specifies the health monitors that the system uses to determine whether it can use this pool for load balancing. Multiple monitors may be specified with the and keyword. The min keyword is used to specify the minimum number of monitors that must succeed for this pool to be declared up. The require keyword is used to specify the minimum number of probes that must succeed for this server to be declared up and the number of probes that should be used. The default value is none.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition within which the component resides.

qos-hit-ratio

Assigns a weight to the Hit Ratio performance factor for the Quality of Service dynamic load balancing mode. The default value is 5.

qos-hops

Assigns a weight to the Hops performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-kilobytes-second

Assigns a weight to the Kilobytes per Second performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 3.

qos-lcs

Assigns a weight to the Link Capacity performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 30.

qos-packet-rate

Assigns a weight to the Packet Rate performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 1.

qos-rtt

Assigns a weight to the Round Trip Time performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 50.

qos-topology

Assigns a weight to the Topology performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-capacity

Assigns a weight to the Virtual Server performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-score

Assigns a weight to the Virtual Server Score performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

tll Specifies the number of seconds that the IP address, once found, is valid. Once the time-to-live (TTL) expires, the client has to request the IP address resolution again. The valid values are 0 through 4294967295; the default value is 30.

verify-member-availability

Specifies that the system verifies the availability of the members before sending a connection to those resources. The default value is enabled.

SEE ALSO

cli admin-partitions, create, delete, edit, glob, gtm monitor, list, ltm default-node-monitor, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-06-09 gtm pool a(1)

NAME

aaaa - Configures AAAA load balancing pools for the Global Traffic Manager(tm).

MODULE

gtm pool

SYNTAX

Modify the Global Traffic Manager pool aaaa component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create pool aaaa [name]
modify pool aaaa [name]
options:
alternate-mode [drop-packet | fallback-ip | global-availability
| none | packet-rate | ratio | return-to-dns | round-robin
| static-persistence | topology | virtual-server-capacity
| virtual-server-score]
app-service [[string] | none]
description [string]
[disabled | enabled]
dynamic ratio [disabled | enabled]
fallback-ip [ip address]
fallback-mode [completion-rate | cpu | drop-packet | fallback-ip
| fewest-hops | global-availability | kilobytes-per-second
| least-connections | lowest-round-trip-time | none
| packet-rate | quality-of-service | ratio | return-to-dns
| round-robin | static-persistence | topology
| virtual-server-capacity | virtual-server-score]
limit-max-bps [integer]
limit-max-bps-status [disabled | enabled]
limit-max-connections [integer]
limit-max-connections-status [disabled | enabled]
limit-max-pps [integer]
limit-max-pps-status [disabled | enabled]
load-balancing-mode [completion-rate | cpu | drop-packet
| fallback-ip | fewest-hops | global-availability
| kilobytes-per-second | least-connections
| lowest-round-trip-time | packet-rate | quality-of-service
| ratio | return-to-dns | round-robin | static-persistence
| topology | virtual-server-capacity | virtual-server-score]
manual-resume [disabled | enabled]
max-answers-returned [integer]
members none
members
[ add | delete | modify | replace-all-with ] {
[server_name:vs-name] {
options:
app-service [[string] | none]
depends-on none
depends-on
[ add | delete | modify | replace-all-with ] {
[server-name:vs-name]...
}
description [string]
[disabled | enabled]
limit-max-bps [integer]
limit-max-bps-status [disabled | enabled]
limit-max-connections [integer]
limit-max-connections-status [disabled | enabled]
limit-max-pps [integer]
limit-max-pps-status [disabled | enabled]
member-order [integer]
monitor [none | [name] [and [name] ]... ]
monitor min [integer] of { [name]... }
monitor require [integer] from [integer] { [name] }
ratio [integer]
}...
}
}
metadata none
metadata
[ add | delete | modify | replace-all-with ] {
[metadata_name ... ] {
app-service [[string] | none]
persist [ true | false ]
value [ "value content" ]
}
}
}
monitor [none | [name] [and [name] ]... ]
monitor min [integer] of { [name]... }
monitor require [integer] from [integer] { [name] }
qos-hit-ratio [integer]
qos-hops [integer]
qos-kilobytes-second [integer]
qos-lcs [integer]
qos-packet-rate [integer]
qos-rtt [integer]
```

qos-topology [integer]
qos-vs-capacity [integer]
qos-vs-score [integer]
ttl [integer]
verify-member-availability [disabled | enabled]

edit pool aaaa [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

reset-stats pool aaaa
reset-stats pool aaaa [[[name] | [glob] | [regex]] ...]

DISPLAY
list pool aaaa
list pool aaaa [[[name] | [glob] | [regex]] ...]
show running-config pool aaaa
show running-config pool aaaa [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
members vs-name
one-line
partition

show pool aaaa
show pool aaaa [name]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

DELETE
delete pool aaaa [name]

Note: You must remove all references to a pool before you can delete the pool.

DESCRIPTION

You can use the pool component to configure the AAAA pool definitions on the Global Traffic Manager. You use a pool to group member servers together to use a common load balancing algorithm.

EXAMPLES

```
create pool aaaa mypool members add {  
  member myServer:myVs  
}  
monitor http
```

Creates a Global Traffic Manager AAAA pool with one member myServer:myVs using the Round Robin load balancing method, and default HTTP monitor checks for member availability.

```
delete pool aaaa my_pool
```

Deletes the pool my_pool.

```
show pool aaaa
```

Displays statistics for all AAAA pools.

```
list pool aaaa my_pool
```

Displays settings of pool my_pool.

OPTIONS

alternate-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred method is unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option. The default value is round-robin.

The options are:

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fallback-ip

Specifies that the Global Traffic Manager returns the IP address that you specify as an answer to the query.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

none Specifies that the system skips the alternate load balancing mode and immediately tries the load balancing mode specified in the fallback-mode option.

Note that if the value of the fallback-mode option is none, and you have multiple pools configured,

the Global Traffic Manager uses the next available pool.

packet-rate
Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

ratio
Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns
Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin
Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence
Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology
Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity
Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

virtual-server-score
Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

app-service
Specifies the name of the application service to which this pool belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool. Only the application service can modify or delete this pool.

description
User defined description.

[disabled | enabled]
Specifies whether this pool is available for load balancing. The default value is enabled.

dynamic-ratio
Enables or disables a dynamic ratio load balancing algorithm for this pool. This option is applicable only when you also configure the load-balancing-mode option for the pool with one of these dynamic load balancing modes: completion-rate, fewest-hops, kilobytes-per-second, least-connections, lowest-round-trip-times, quality-of-service, virtual-server-capacity, or virtual-server-score.

When this option is disabled (the default), the system uses only the server or virtual server with the best metrics, or highest quality of service (QoS) score, for load balancing. When dynamic-ratio is enabled, the system treats QoS scores as ratios, and it uses each server or virtual server in proportion to the ratio determined by the QoS calculation.

fallback-ip
Specifies the IPv6 address of the server to which the system directs requests in the event that the load balancing methods configured for this pool fail to return a valid virtual server. The default value is ::.

fallback-mode
Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred and alternate modes are unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option, and the alternate mode using the alternate-mode option. The default value is return-to-dns.

The options are:

completion-rate
Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

cpu Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet
Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fallback-ip
Specifies that the Global Traffic Manager returns the IP address that you specify as an answer to the query.

fewest-hops
Specifies that the Global Traffic Manager distributes connection requests to the virtual server in

the data center that has the fewest router hops from the Local DNS.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

none Specifies that there is no fallback mode. If the system cannot use the preferred or alternate load balancing modes, it uses the next pool to resolve the request. If there are no more pools available, the result is the same as when the value for the fallback-mode option is return-to-dns.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

limit-max-bps

Specifies the maximum allowable data throughput rate, in bits per second, for the virtual servers in the pool. If the network traffic volume exceeds this value, the system marks the pool as unavailable. The default value is 0 (zero).

limit-max-bps-status

Enables or disables the limit-max-bps option for this pool. The default value is disabled.

limit-max-connections

Specifies the number of current connections allowed for the virtual servers in the pool. If the current connections exceed this value, the system marks the pool as unavailable. The default value is 0 (zero).

limit-max-connections-status

Enables or disables the limit-max-connections option for this pool. The default value is disabled.

limit-max-pps

Specifies the maximum allowable data transfer rate, in packets per second, for the virtual servers in the pool. If the network traffic volume exceeds this value, the system marks the pool as unavailable. The default value is 0 (zero).

limit-max-pps-status

Enables or disables the limit-max-pps option for this pool. The default value is disabled.

load-balancing-mode

Specifies the preferred load balancing mode that the system uses to load balance name resolution requests among the members of this pool. The default value is round-robin.

The options are:

completion-rate

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

`cpu` Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fallback-ip

Specifies that the Global Traffic Manager returns the IP address that you specify as an answer to the query.

fewest-hops

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

manual-resume

Enables or disables the manual resume function for this pool. If you leave this option disabled (the default), then a member of this pool automatically becomes available for load balancing when its status changes from down to up. When the manual-resume option is enabled, if the status of a member of this pool changes from up to down, the pool member remains disabled indefinitely until you manually re-enable it.

max-answers-returned

Specifies the maximum number of available virtual servers that the system lists in a response. The default value is 1.

members

Specifies the vs-name of the pool members. The default value is none.

You can also use the following options with pool members:

app-service

Specifies the name of the application service to which this pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool member. Only the application service can modify or delete this pool member.

depends-on

Specifies the name of the virtual server on which this pool member depends.

description

User defined description.

[enabled | disabled]

Specifies whether this pool member is available for load balancing. The default value is enabled.

limit-max-bps

Specifies the maximum allowable data throughput rate, in bits per second, for the pool member. If the network traffic volume exceeds this value, the system marks the pool member as unavailable.

limit-max-bps-status

Enables or disables the limit-max-bps option for this pool member. The default value is disabled.

limit-max-connections

Specifies the number of current connections allowed for this pool member. If the current connections exceed this value, the system marks this pool member as unavailable.

limit-max-connections-status

Enables or disables the limit-max-connection option for this pool member. The default value is disabled.

limit-max-pps

Specifies the maximum allowable data transfer rate, in packets per second, for this pool member. If the network traffic volume exceeds this value, the system marks this pool member as unavailable.

limit-max-pps-status

Enables or disables the limit-max-pps option for this pool member. The default value is disabled.

member-order

Specifies the order number of the pool member. The system uses this number with load balancing methods that involve prioritizing pool members, such as the Ratio load balancing method.

monitor

Specifies the health monitors that the system uses to determine whether it can use this pool member for load balancing. Multiple monitors may be specified with the and keyword. The min keyword is used to specify the minimum number of monitors that must succeed for this pool member to be declared up. The require keyword is used to specify the minimum number of probes that must succeed for this server to be declared up and the number of probes that should be used. The default value is none.

ratio

Specifies the weight of the pool member for load balancing purposes.

vs-name

Displays the name of the corresponding virtual server.

metadata

Associates user defined data, each of which has name and value pair and persistence. Persistent(default) means the data will be saved into config file.

monitor

Specifies the health monitors that the system uses to determine whether it can use this pool for load balancing. Multiple monitors may be specified with the and keyword. The min keyword is used to specify the minimum number of monitors that must succeed for this pool to be declared up. The require keyword is used to specify the minimum number of probes that must succeed for this server to be declared up and the number of probes that should be used. The default value is none.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition within which the component resides.

qos-hit-ratio

Assigns a weight to the Hit Ratio performance factor for the Quality of Service dynamic load balancing mode. The default value is 5.

qos-hops

Assigns a weight to the Hops performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-kilobytes-second

Assigns a weight to the Kilobytes per Second performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 3.

qos-lcs

Assigns a weight to the Link Capacity performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 30.

qos-packet-rate

Assigns a weight to the Packet Rate performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 1.

qos-rtt

Assigns a weight to the Round Trip Time performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 50.

qos-topology

Assigns a weight to the Topology performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-capacity

Assigns a weight to the Virtual Server performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-score

Assigns a weight to the Virtual Server Score performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ttl Specifies the number of seconds that the IP address, once found, is valid. Once the time-to-live (TTL) expires, the client has to request the IP address resolution again. The valid values are 0 through 4294967295; the default value is 30.

verify-member-availability

Specifies that the system verifies the availability of the members before sending a connection to those resources. The default value is enabled.

SEE ALSO

cli admin-partitions, create, delete, edit, glob, gtm monitor, list, ltm default-node-monitor, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-06-10 gtm pool aaaa(1)

gtm pool cname

NAME

cname - Configures CNAME load balancing pools for the Global Traffic Manager(tm).

MODULE

gtm pool

SYNTAX

Modify the Global Traffic Manager pool cname component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

create pool cname [name]

modify pool cname [name]

options:

alternate-mode [drop-packet | global-availability
| none | packet-rate | ratio | return-to-dns | round-robin
| static-persistence | topology | virtual-server-capacity
| virtual-server-score]

app-service [[string] | none]

description [string]

[disabled | enabled]

dynamic ratio [disabled | enabled]

fallback-mode [completion-rate | cpu | drop-packet

```

| fewest-hops | global-availability | kilobytes-per-second
| least-connections | lowest-round-trip-time | none
| packet-rate | quality-of-service | ratio | return-to-dns
| round-robin | static-persistence | topology
| virtual-server-capacity | virtual-server-score]
load-balancing-mode [completion-rate | cpu | drop-packet
| fewest-hops | global-availability
| kilobytes-per-second | least-connections
| lowest-round-trip-time | packet-rate | quality-of-service
| ratio | return-to-dns | round-robin | static-persistence
| topology | virtual-server-capacity | virtual-server-score]
manual-resume [disabled | enabled]
members none
members
[ add | delete | modify | replace-all-with ] {
  [member-dname] {
options:
app-service [[string] | none]
description [string]
[disabled | enabled]
member-order [integer]
ratio [integer]
static-target [yes | no]
}...
}
metadata none
metadata
[add | delete | modify | replace-all-with] {
  [metadata_name ... ] {
app-service [[string] | none]
persist [ true | false ]
value [ "value content" ]
}
}
qos-hit-ratio [integer]
qos-hops [integer]
qos-kilobytes-second [integer]
qos-lcs [integer]
qos-packet-rate [integer]
qos-rtt [integer]
qos-topology [integer]
qos-vs-capacity [integer]
qos-vs-score [integer]
ttl [integer]
verify-member-availability [disabled | enabled]

```

```
edit pool cname [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
all-properties
non-default-properties
```

```
reset-stats pool cname
reset-stats pool cname [ [ [name] | [glob] | [regex] ] ... ]
```

```

DISPLAY
list pool cname
list pool cname [ [ [name] | [glob] | [regex] ] ... ]
show running-config pool cname
show running-config pool cname [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties
members member-dname
one-line
partition

```

```

show pool cname
show pool cname [name]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

```

```

DELETE
delete pool cname [name]

```

Note: You must remove all references to a pool before you can delete the pool.

DESCRIPTION

You can use the pool component to configure the CNAME pool definitions on the Global Traffic Manager. You use a pool to group member names together to use a common load balancing algorithm.

EXAMPLES

```

create pool cname mypool members add {
  canonical.mydomain.com
}

```

Creates a Global Traffic Manager CNAME pool with one member canonical.mydomain.com using the Round Robin load balancing method.

delete pool cname my_pool

Deletes the pool my_pool.

show pool cname

Displays statistics for all CNAME pools.

list pool cname my_pool

Displays settings of pool my_pool.

OPTIONS

alternate-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred method is unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option. The default value is round-robin.

The options are:

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

none Specifies that the system skips the alternate load balancing mode and immediately tries the load balancing mode specified in the fallback-mode option.

Note that if the value of the fallback-mode option is none, and you have multiple pools configured, the Global Traffic Manager uses the next available pool.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

app-service

Specifies the name of the application service to which this pool belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool. Only the application service can modify or delete this pool.

description

User defined description.

[disabled | enabled]

Specifies whether this pool is available for load balancing. The default value is enabled.

dynamic-ratio

Enables or disables a dynamic ratio load balancing algorithm for this pool. This option is applicable only when you also configure the load-balancing-mode option for the pool with one of these dynamic load balancing modes: completion-rate, fewest-hops, kilobytes-per-second, least-connections, lowest-round-trip-times, quality-of-service, virtual-server-capacity, or virtual-server-score.

When this option is disabled (the default), the system uses only the server or virtual server with the

best metrics, or highest quality of service (QoS) score, for load balancing. When dynamic-ratio is enabled, the system treats QoS scores as ratios, and it uses each server or virtual server in proportion to the ratio determined by the QoS calculation.

fallback-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred and alternate modes are unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option, and the alternate mode using the alternate-mode option. The default value is return-to-dns.

The options are:

completion-rate

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

`cpu` Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fewest-hops

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

`none` Specifies that there is no fallback mode. If the system cannot use the preferred or alternate load balancing modes, it uses the next pool to resolve the request. If there are no more pools available, the result is the same as when the value for the fallback-mode option is return-to-dns.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`load-balancing-mode`
Specifies the preferred load balancing mode that the system uses to load balance name resolution requests among the members of this pool. The default value is `round-robin`.

The options are:

`completion-rate`
Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

`cpu` Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

`drop-packet`
Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

`fewest-hops`
Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

`global-availability`
Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

`kilobytes-per-second`
Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

`least-connections`
Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

`lowest-round-trip-time`
Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

`packet-rate`
Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

`quality-of-service`
Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

`ratio`
Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

`return-to-dns`
Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

`round-robin`
Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

`static-persistence`
Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

`topology`
Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

`virtual-server-capacity`
Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

`virtual-server-score`
Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`manual-resume`
Enables or disables the manual resume function for this pool. If you leave this option disabled (the default), then a member of this pool automatically becomes available for load balancing when its status changes from down to up. When the manual-resume option is enabled, if the status of a member of this pool changes from up to down, the pool member remains disabled indefinitely until you manually re-enable it.

`members`

Specifies the member-dname of the pool members. The default value is none.

You can also use the following options with pool members:

app-service

Specifies the name of the application service to which this pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool member. Only the application service can modify or delete this pool member.

description

User defined description.

[enabled | disabled]

Specifies whether this pool member is available for load balancing. The default value is enabled.

member-order

Specifies the order number of the pool member. The system uses this number with load balancing methods that involve prioritizing pool members, such as the Ratio load balancing method.

ratio

Specifies the weight of the pool member for load balancing purposes.

static-target

Specifies the member's name specifies a static DNAME rather than a name linked to a Wide IP defined on the system. This may be required if the target DNAME is not owned by the organization or configured on GTM. One side-effect of using a static target is that the member is always considered available for load balancing. The default value is no.

member-dname

Displays the member's DNAME.

metadata

Associates user defined data, each of which has name and value pair and persistence. Persistent(default) means the data will be saved into config file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition within which the component resides.

qos-hit-ratio

Assigns a weight to the Hit Ratio performance factor for the Quality of Service dynamic load balancing mode. The default value is 5.

qos-hops

Assigns a weight to the Hops performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-kilobytes-second

Assigns a weight to the Kilobytes per Second performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 3.

qos-lcs

Assigns a weight to the Link Capacity performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 30.

qos-packet-rate

Assigns a weight to the Packet Rate performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 1.

qos-rtt

Assigns a weight to the Round Trip Time performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 50.

qos-topology

Assigns a weight to the Topology performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-capacity

Assigns a weight to the Virtual Server performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-score

Assigns a weight to the Virtual Server Score performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

tll Specifies the number of seconds that the answer, once found, is valid. Once the time-to-live (TTL) expires, the client has to request name resolution again. The valid values are 0 through 4294967295; the default value is 30.

verify-member-availability

Specifies that the system verifies the availability of the members before sending a connection to those resources. The default value is enabled.

SEE ALSO

cli admin-partitions, create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-04-28 gtm pool cname(1)

gtm pool mx

NAME

mx - Configures MX load balancing pools for the Global Traffic Manager(tm).

MODULE

gtm pool

SYNTAX

Modify the Global Traffic Manager pool mx component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create pool mx [name]
modify pool mx [name]
options:
  alternate-mode [drop-packet | global-availability
    | none | packet-rate | ratio | return-to-dns | round-robin
    | static-persistence | topology | virtual-server-capacity
    | virtual-server-score]
  app-service [[string] | none]
  description [string]
  [disabled | enabled]
  dynamic ratio [disabled | enabled]
  fallback-mode [completion-rate | cpu | drop-packet
    | fewest-hops | global-availability | kilobytes-per-second
    | least-connections | lowest-round-trip-time | none
    | packet-rate | quality-of-service | ratio | return-to-dns
    | round-robin | static-persistence | topology
    | virtual-server-capacity | virtual-server-score]
  load-balancing-mode [completion-rate | cpu | drop-packet
    | fewest-hops | global-availability
    | kilobytes-per-second | least-connections
    | lowest-round-trip-time | packet-rate | quality-of-service
    | ratio | return-to-dns | round-robin | static-persistence
    | topology | virtual-server-capacity | virtual-server-score]
  manual-resume [disabled | enabled]
  max-answers-returned [integer]
  members none
  members
    [ add | delete | modify | replace-all-with ] {
      [member-dname] {
options:
  app-service [[string] | none]
  description [string]
  [disabled | enabled]
  member-order [integer]
  priority [integer]
  ratio [integer]
  }...
    }
  metadata none
  metadata
    [add | delete | modify | replace-all-with] {
      [metadata_name ... ] {
  app-service [[string] | none]
  persist [ true | false ]
  value [ "value content" ]
    }
  }
  qos-hit-ratio [integer]
  qos-hops [integer]
  qos-kilobytes-second [integer]
  qos-lcs [integer]
```

```
qos-packet-rate [integer]
qos-rtt [integer]
qos-topology [integer]
qos-vs-capacity [integer]
qos-vs-score [integer]
ttl [integer]
verify-member-availability [disabled | enabled]
```

```
edit pool mx [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

```
reset-stats pool mx
```

```
reset-stats pool mx [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list pool mx
```

```
list pool mx [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config pool mx
```

```
show running-config pool mx [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
members member-dname
one-line
partition
```

```
show pool mx
```

```
show pool mx [name]
```

options:

```
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt
```

DELETE

```
delete pool mx [name]
```

Note: You must remove all references to a pool before you can delete the pool.

DESCRIPTION

You can use the pool component to configure the MX pool definitions on the Global Traffic Manager. You use a pool to group member names together to use a common load balancing algorithm.

EXAMPLES

```
create pool mx mypool members add {
  mail1.mydomain.com {
    priority 20
  }
}
```

Creates a Global Traffic Manager MX pool with one member mail1.mydomain.com with priority 20 using the Round Robin load balancing method.

```
delete pool mx my_pool
```

Deletes the pool my_pool.

```
show pool mx
```

Displays statistics for all MX pools.

```
list pool mx my_pool
```

Displays settings of pool my_pool.

OPTIONS

alternate-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred method is unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option. The default value is round-robin.

The options are:

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

none Specifies that the system skips the alternate load balancing mode and immediately tries the load balancing mode specified in the fallback-mode option.

Note that if the value of the fallback-mode option is none, and you have multiple pools configured, the Global Traffic Manager uses the next available pool.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

app-service

Specifies the name of the application service to which this pool belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool. Only the application service can modify or delete this pool.

description

User defined description.

[disabled | enabled]

Specifies whether this pool is available for load balancing. The default value is enabled.

dynamic-ratio

Enables or disables a dynamic ratio load balancing algorithm for this pool. This option is applicable only when you also configure the load-balancing-mode option for the pool with one of these dynamic load balancing modes: completion-rate, fewest-hops, kilobytes-per-second, least-connections, lowest-round-trip-times, quality-of-service, virtual-server-capacity, or virtual-server-score.

When this option is disabled (the default), the system uses only the server or virtual server with the best metrics, or highest quality of service (QoS) score, for load balancing. When dynamic-ratio is enabled, the system treats QoS scores as ratios, and it uses each server or virtual server in proportion to the ratio determined by the QoS calculation.

fallback-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred and alternate modes are unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option, and the alternate mode using the alternate-mode option. The default value is return-to-dns.

The options are:

completion-rate

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

cpu Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fewest-hops

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

none Specifies that there is no fallback mode. If the system cannot use the preferred or alternate load balancing modes, it uses the next pool to resolve the request. If there are no more pools available, the result is the same as when the value for the fallback-mode option is return-to-dns.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

load-balancing-mode

Specifies the preferred load balancing mode that the system uses to load balance name resolution requests among the members of this pool. The default value is round-robin.

The options are:

completion-rate

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

cpu Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fewest-hops

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time
Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

packet-rate
Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service
Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio
Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns
Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin
Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence
Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology
Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity
Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score
Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

manual-resume
Enables or disables the manual resume function for this pool. If you leave this option disabled (the default), then a member of this pool automatically becomes available for load balancing when its status changes from down to up. When the manual-resume option is enabled, if the status of a member of this pool changes from up to down, the pool member remains disabled indefinitely until you manually re-enable it.

max-answers-returned
Specifies the maximum number of available virtual servers that the system lists in a response. The default value is 1.

members
Specifies the member-dname of the pool members. The default value is none.

You can also use the following options with pool members:

app-service
Specifies the name of the application service to which this pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool member. Only the application service can modify or delete this pool member.

description
User defined description.

[enabled | disabled]
Specifies whether this pool member is available for load balancing. The default value is enabled.

member-order
Specifies the order number of the pool member. The system uses this number with load balancing methods that involve prioritizing pool members, such as the Ratio load balancing method.

priority
Specifies the response resource record's priority RDATA field value when this member is picked. The default value is 10.

ratio
Specifies the weight of the pool member for load balancing purposes.

member-dname
Displays the member's DNAME.

metadata
Associates user defined data, each of which has name and value pair and persistence. Persistent(default)

means the data will be saved into config file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition within which the component resides.

qos-hit-ratio

Assigns a weight to the Hit Ratio performance factor for the Quality of Service dynamic load balancing mode. The default value is 5.

qos-hops

Assigns a weight to the Hops performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-kilobytes-second

Assigns a weight to the Kilobytes per Second performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 3.

qos-lcs

Assigns a weight to the Link Capacity performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 30.

qos-packet-rate

Assigns a weight to the Packet Rate performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 1.

qos-rtt

Assigns a weight to the Round Trip Time performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 50.

qos-topology

Assigns a weight to the Topology performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-capacity

Assigns a weight to the Virtual Server performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-score

Assigns a weight to the Virtual Server Score performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ttd Specifies the number of seconds that the answer, once found, is valid. Once the time-to-live (TTL)

expires, the client has to request name resolution again. The valid values are 0 through 4294967295; the default value is 30.

verify-member-availability

Specifies that the system verifies the availability of the members before sending a connection to those resources. The default value is enabled.

SEE ALSO

cli admin-partitions, create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-04-28 gtm pool mx(1)

gtm pool naptr

NAME

naptr - Configures NAPTR load balancing pools for the Global Traffic Manager(tm).

MODULE

gtm pool

SYNTAX

Modify the Global Traffic Manager pool naptr component within the gtm module using the syntax shown in the following sections.

```

CREATE/MODIFY
create pool naptr [name]
modify pool naptr [name]
options:
alternate-mode [drop-packet | global-availability
| none | packet-rate | ratio | return-to-dns | round-robin
| static-persistence | topology | virtual-server-capacity
| virtual-server-score]
app-service [[string] | none]
description [string]
[disabled | enabled]
dynamic ratio [disabled | enabled]
fallback-mode [completion-rate | cpu | drop-packet
| fewest-hops | global-availability | kilobytes-per-second
| least-connections | lowest-round-trip-time | none
| packet-rate | quality-of-service | ratio | return-to-dns
| round-robin | static-persistence | topology
| virtual-server-capacity | virtual-server-score]
load-balancing-mode [completion-rate | cpu | drop-packet
| fewest-hops | global-availability
| kilobytes-per-second | least-connections
| lowest-round-trip-time | packet-rate | quality-of-service
| ratio | return-to-dns | round-robin | static-persistence
| topology | virtual-server-capacity | virtual-server-score]
manual-resume [disabled | enabled]
max-answers-returned [integer]
members none
members
[ add | delete | modify | replace-all-with ] {
[member-dname] {
options:
app-service [[string] | none]
description [string]
[disabled | enabled]
flags [string]
member-order [integer]
order [integer]
preference [integer]
ratio [integer]
service [string]
}...
}
metadata none
metadata
[add | delete | modify | replace-all-with] {
[metadata_name ... ] {
app-service [[string] | none]
persist [ true | false ]
value [ "value content" ]
}
}
qos-hit-ratio [integer]
qos-hops [integer]
qos-kilobytes-second [integer]
qos-lcs [integer]
qos-packet-rate [integer]
qos-rtt [integer]
qos-topology [integer]
qos-vs-capacity [integer]
qos-vs-score [integer]
ttl [integer]
verify-member-availability [disabled | enabled]

edit pool naptr [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties

reset-stats pool naptr
reset-stats pool naptr [ [ [name] | [glob] | [regex] ] ... ]

DISPLAY
list pool naptr
list pool naptr [ [ [name] | [glob] | [regex] ] ... ]
show running-config pool naptr
show running-config pool naptr [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties
members member-dname
one-line
partition

show pool naptr
show pool naptr [name]
options:

```

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

DELETE
delete pool naptr [name]

Note: You must remove all references to a pool before you can delete the pool.

DESCRIPTION

You can use the pool component to configure the NAPTR pool definitions on the Global Traffic Manager. You use a pool to group member names together to use a common load balancing algorithm.

EXAMPLES

```
create pool naptr mypool members add {
  _sip_udp.a.mydomain.com {
    flags "s"
    order 20
    preference 100
    service "sip+d2u"
  }
  _sip_udp.b.mydomain.com {
    flags "s"
    order 20
    preference 110
    service "sip+d2u"
  }
}
```

Creates a Global Traffic Manager NAPTR pool with two members using the Round Robin load balancing method.

```
delete pool naptr my_pool
```

Deletes the pool my_pool.

```
show pool naptr
```

Displays statistics for all NAPTR pools.

```
list pool naptr my_pool
```

Displays settings of pool my_pool.

OPTIONS

alternate-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred method is unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option. The default value is round-robin.

The options are:

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

none Specifies that the system skips the alternate load balancing mode and immediately tries the load balancing mode specified in the fallback-mode option.

Note that if the value of the fallback-mode option is none, and you have multiple pools configured, the Global Traffic Manager uses the next available pool.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute

connection requests.

`virtual-server-capacity`

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`virtual-server-score`

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`app-service`

Specifies the name of the application service to which this pool belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete this pool. Only the application service can modify or delete this pool.

`description`

User defined description.

`[disabled | enabled]`

Specifies whether this pool is available for load balancing. The default value is enabled.

`dynamic-ratio`

Enables or disables a dynamic ratio load balancing algorithm for this pool. This option is applicable only when you also configure the `load-balancing-mode` option for the pool with one of these dynamic load balancing modes: `completion-rate`, `fewest-hops`, `kilobytes-per-second`, `least-connections`, `lowest-round-trip-times`, `quality-of-service`, `virtual-server-capacity`, or `virtual-server-score`.

When this option is disabled (the default), the system uses only the server or virtual server with the best metrics, or highest quality of service (QoS) score, for load balancing. When `dynamic-ratio` is enabled, the system treats QoS scores as ratios, and it uses each server or virtual server in proportion to the ratio determined by the QoS calculation.

`fallback-mode`

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred and alternate modes are unsuccessful in picking a pool. You set the preferred mode using the `load-balancing-mode` option, and the alternate mode using the `alternate-mode` option. The default value is `return-to-dns`.

The options are:

`completion-rate`

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

`cpu` Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

`drop-packet`

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

`fewest-hops`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

`global-availability`

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

`kilobytes-per-second`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

`least-connections`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

`lowest-round-trip-time`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

`none` Specifies that there is no fallback mode. If the system cannot use the preferred or alternate load balancing modes, it uses the next pool to resolve the request. If there are no more pools available, the result is the same as when the value for the `fallback-mode` option is `return-to-dns`.

`packet-rate`

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

`quality-of-service`

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

`ratio`

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual

servers using a weighted Round Robin load balancing method.

`return-to-dns`

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

`round-robin`

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

`static-persistence`

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

`topology`

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

`virtual-server-capacity`

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

`virtual-server-score`

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`glob` Displays the items that match the `glob` expression. See help `glob` for a description of `glob` expression syntax.

`load-balancing-mode`

Specifies the preferred load balancing mode that the system uses to load balance name resolution requests among the members of this pool. The default value is `round-robin`.

The options are:

`completion-rate`

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

`cpu` Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

`drop-packet`

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

`fewest-hops`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

`global-availability`

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

`kilobytes-per-second`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

`least-connections`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

`lowest-round-trip-time`

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

`packet-rate`

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

`quality-of-service`

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

`ratio`

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

`return-to-dns`

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

`round-robin`

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

`static-persistence`
Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

`topology`
Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

`virtual-server-capacity`
Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

`virtual-server-score`
Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`manual-resume`
Enables or disables the manual resume function for this pool. If you leave this option disabled (the default), then a member of this pool automatically becomes available for load balancing when its status changes from down to up. When the manual-resume option is enabled, if the status of a member of this pool changes from up to down, the pool member remains disabled indefinitely until you manually re-enable it.

`max-answers-returned`
Specifies the maximum number of available virtual servers that the system lists in a response. The default value is 1.

`members`
Specifies the member-dname of the pool members. The default value is none.

You can also use the following options with pool members:

`app-service`
Specifies the name of the application service to which this pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool member. Only the application service can modify or delete this pool member.

`description`
User defined description.

[enabled | disabled]
Specifies whether this pool member is available for load balancing. The default value is enabled.

`flags`
Specifies the response resource record's flags RDATA field value when this member is picked. Either 'a' or 's' may be specified for this attribute as GTM supports only these record types.

`member-order`
Specifies the order number of the pool member. The system uses this number with load balancing methods that involve prioritizing pool members, such as the Ratio load balancing method.

`order`
Specifies the response resource record's order RDATA field value when this member is picked. The default value is 10.

`preference`
Specifies the response resource record's preference RDATA field value when this member is picked. The default value is 10.

`ratio`
Specifies the weight of the pool member for load balancing purposes.

`service`
Specifies the response resource record's service RDATA field value when this member is picked.

`member-dname`
Displays the member's DNAME.

`metadata`
Associates user defined data, each of which has name and value pair and persistence. Persistent(default) means the data will be saved into config file.

`name` Specifies a unique name for the component. This option is required for the commands create and modify.

`partition`
Displays the partition within which the component resides.

`qos-hit-ratio`
Assigns a weight to the Hit Ratio performance factor for the Quality of Service dynamic load balancing mode. The default value is 5.

`qos-hops`

Assigns a weight to the Hops performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-kilobytes-second

Assigns a weight to the Kilobytes per Second performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 3.

qos-lcs

Assigns a weight to the Link Capacity performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 30.

qos-packet-rate

Assigns a weight to the Packet Rate performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 1.

qos-rtt

Assigns a weight to the Round Trip Time performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 50.

qos-topology

Assigns a weight to the Topology performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-capacity

Assigns a weight to the Virtual Server performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-score

Assigns a weight to the Virtual Server Score performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ttl Specifies the number of seconds that the answer, once found, is valid. Once the time-to-live (TTL) expires, the client has to request name resolution again. The valid values are 0 through 4294967295; the default value is 30.

verify-member-availability

Specifies that the system verifies the availability of the members before sending a connection to those resources. The default value is enabled.

SEE ALSO

cli admin-partitions, create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-04-28 gtm pool naptr(1)

gtm pool srv

NAME

srv - Configures SRV load balancing pools for the Global Traffic Manager(tm).

MODULE

gtm pool

SYNTAX

Modify the Global Traffic Manager pool srv component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

create pool srv [name]
modify pool srv [name]

options:

alternate-mode [drop-packet | global-availability
| none | packet-rate | ratio | return-to-dns | round-robin
| static-persistence | topology | virtual-server-capacity
| virtual-server-score]

app-service [[string] | none]

description [string]

[disabled | enabled]

```

dynamic ratio [disabled | enabled]
fallback-mode [completion-rate | cpu | drop-packet
| fewest-hops | global-availability | kilobytes-per-second
| least-connections | lowest-round-trip-time | none
| packet-rate | quality-of-service | ratio | return-to-dns
| round-robin | static-persistence | topology
| virtual-server-capacity | virtual-server-score]
load-balancing-mode [completion-rate | cpu | drop-packet
| fewest-hops | global-availability
| kilobytes-per-second | least-connections
| lowest-round-trip-time | packet-rate | quality-of-service
| ratio | return-to-dns | round-robin | static-persistence
| topology | virtual-server-capacity | virtual-server-score]
manual-resume [disabled | enabled]
max-answers-returned [integer]
members none
members
[ add | delete | modify | replace-all-with ] {
[member-dname] {
options:
app-service [[string] | none]
description [string]
[disabled | enabled]
member-order [integer]
port [integer]
priority [integer]
ratio [integer]
weight [integer]
}...
}
metadata none
metadata
[add | delete | modify | replace-all-with] {
[metadata_name ... ] {
app-service [[string] | none]
persist [ true | false ]
value [ "value content" ]
}
}
qos-hit-ratio [integer]
qos-hops [integer]
qos-kilobytes-second [integer]
qos-lcs [integer]
qos-packet-rate [integer]
qos-rtt [integer]
qos-topology [integer]
qos-vs-capacity [integer]
qos-vs-score [integer]
ttl [integer]
verify-member-availability [disabled | enabled]

```

```

edit pool srv [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties

```

```

reset-stats pool srv
reset-stats pool srv [ [ [name] | [glob] | [regex] ] ... ]

```

```

DISPLAY
list pool srv
list pool srv [ [ [name] | [glob] | [regex] ] ... ]
show running-config pool srv
show running-config pool srv [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties
members member-dname
one-line
partition

```

```

show pool srv
show pool srv [name]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

```

```

DELETE
delete pool srv [name]

```

Note: You must remove all references to a pool before you can delete the pool.

DESCRIPTION

You can use the pool component to configure the SRV pool definitions on the Global Traffic Manager. You use a pool to group member names together to use a common load balancing algorithm.

EXAMPLES

```
create pool srv mypool members add {
  imap.mydomain.com {
    port 143
    priority 20
  }
}
```

Creates a Global Traffic Manager SRV pool with one member `imap.mydomain.com` with port 143 and priority 20 using the Round Robin load balancing method.

```
delete pool srv my_pool
```

Deletes the pool `my_pool`.

```
show pool srv
```

Displays statistics for all SRV pools.

```
list pool srv my_pool
```

Displays settings of pool `my_pool`.

OPTIONS

`alternate-mode`

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred method is unsuccessful in picking a pool. You set the preferred mode using the `load-balancing-mode` option. The default value is `round-robin`.

The options are:

`drop-packet`

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

`global-availability`

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

`none` Specifies that the system skips the alternate load balancing mode and immediately tries the load balancing mode specified in the `fallback-mode` option.

Note that if the value of the `fallback-mode` option is `none`, and you have multiple pools configured, the Global Traffic Manager uses the next available pool.

`packet-rate`

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

`ratio`

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

`return-to-dns`

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

`round-robin`

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

`static-persistence`

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

`topology`

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

`virtual-server-capacity`

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`virtual-server-score`

Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

`app-service`

Specifies the name of the application service to which this pool belongs. The default value is `none`.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete this pool. Only the application service can modify or delete this pool.

`description`

User defined description.

[disabled | enabled]

Specifies whether this pool is available for load balancing. The default value is `enabled`.

dynamic-ratio

Enables or disables a dynamic ratio load balancing algorithm for this pool. This option is applicable only when you also configure the load-balancing-mode option for the pool with one of these dynamic load balancing modes: completion-rate, fewest-hops, kilobytes-per-second, least-connections, lowest-round-trip-times, quality-of-service, virtual-server-capacity, or virtual-server-score.

When this option is disabled (the default), the system uses only the server or virtual server with the best metrics, or highest quality of service (QoS) score, for load balancing. When dynamic-ratio is enabled, the system treats QoS scores as ratios, and it uses each server or virtual server in proportion to the ratio determined by the QoS calculation.

fallback-mode

Specifies the load balancing mode that the system uses to load balance name resolution requests among the members of this pool, if the preferred and alternate modes are unsuccessful in picking a pool. You set the preferred mode using the load-balancing-mode option, and the alternate mode using the alternate-mode option. The default value is return-to-dns.

The options are:

completion-rate

Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

`cpu` Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet

Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fewest-hops

Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

global-availability

Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second

Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections

Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time

Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

`none` Specifies that there is no fallback mode. If the system cannot use the preferred or alternate load balancing modes, it uses the next pool to resolve the request. If there are no more pools available, the result is the same as when the value for the fallback-mode option is return-to-dns.

packet-rate

Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service

Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio

Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns

Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin

Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence

Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology

Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity

Specifies that the Global Traffic Manager distributes connection requests by creating a list of the

virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score
Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

load-balancing-mode
Specifies the preferred load balancing mode that the system uses to load balance name resolution requests among the members of this pool. The default value is round-robin.

The options are:

completion-rate
Specifies that the Global Traffic Manager selects the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client LDNS.

cpu Specifies that the Global Traffic Manager selects the virtual server that currently has the most CPU processing time available to handle name resolution requests.

drop-packet
Specifies that the Global Traffic Manager does nothing with the packet, and simply drops the request.

fewest-hops
Specifies that the Global Traffic Manager distributes connection requests to the virtual server in the data center that has the fewest router hops from the Local DNS.

global-availability
Specifies that the Global Traffic Manager distributes connection requests to virtual servers included in the pool in the order in which they are listed.

kilobytes-per-second
Specifies that the Global Traffic Manager distributes connection requests to the virtual server that is currently processing the fewest number of kilobytes per second.

least-connections
Specifies that the Global Traffic Manager distributes connection requests to the virtual server on the Local Traffic Manager that currently hosts the fewest connections.

lowest-round-trip-time
Specifies that the Global Traffic Manager distributes connection requests to the virtual server with the fastest measured round trip time between a data center and a client LDNS.

packet-rate
Specifies that the Global Traffic Manager assigns connection requests to the virtual server that is currently processing the fewest number of packets per second.

quality-of-service
Specifies that the Global Traffic Manager distributes connection requests using current performance information to calculate an overall score for each virtual server, and then distributes connections to the virtual servers based on these scores.

ratio
Specifies that the Global Traffic Manager distributes connection requests among a pool of virtual servers using a weighted Round Robin load balancing method.

return-to-dns
Specifies that the Global Traffic Manager immediately returns connection requests to the Local DNS for resolution.

round-robin
Specifies that the Global Traffic Manager distributes connection requests in a circular and sequential pattern among the virtual servers in a pool.

static-persistence
Specifies that the Global Traffic Manager consistently maps an LDNS IP address to the same available virtual server for the duration of a session.

topology
Specifies that the Global Traffic Manager uses proximity-based load balancing to distribute connection requests.

virtual-server-capacity
Specifies that the Global Traffic Manager distributes connection requests by creating a list of the virtual servers, weighted by capacity, then picks one of the virtual servers from the list. The virtual servers with the greatest capacity are picked most often, but over time all virtual servers are returned.

virtual-server-score
Specifies that the Global Traffic Manager assigns connection requests to virtual servers based on a user-defined ranking system.

manual-resume

Enables or disables the manual resume function for this pool. If you leave this option disabled (the default), then a member of this pool automatically becomes available for load balancing when its status changes from down to up. When the manual-resume option is enabled, if the status of a member of this pool changes from up to down, the pool member remains disabled indefinitely until you manually re-enable it.

max-answers-returned

Specifies the maximum number of available virtual servers that the system lists in a response. The default value is 1.

members

Specifies the member-dname of the pool members. The default value is none.

You can also use the following options with pool members:

app-service

Specifies the name of the application service to which this pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this pool member. Only the application service can modify or delete this pool member.

description

User defined description.

[enabled | disabled]

Specifies whether this pool member is available for load balancing. The default value is enabled.

member-order

Specifies the order number of the pool member. The system uses this number with load balancing methods that involve prioritizing pool members, such as the Ratio load balancing method.

port Specifies the response resource record's port RDATA field value when this member is picked.

priority

Specifies the response resource record's priority RDATA field value when this member is picked. The default value is 10.

ratio

Specifies the relative weight of the pool member for load balancing purposes.

weight

Specifies the response resource record's weight RDATA field value when this member is picked. The default value is 10.

member-dname

Displays the member's DNAME.

metadata

Associates user defined data, each of which has name and value pair and persistence. Persistent(default) means the data will be saved into config file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition within which the component resides.

qos-hit-ratio

Assigns a weight to the Hit Ratio performance factor for the Quality of Service dynamic load balancing mode. The default value is 5.

qos-hops

Assigns a weight to the Hops performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-kilobytes-second

Assigns a weight to the Kilobytes per Second performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 3.

qos-lcs

Assigns a weight to the Link Capacity performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 30.

qos-packet-rate

Assigns a weight to the Packet Rate performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 1.

qos-rtt

Assigns a weight to the Round Trip Time performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 50.

qos-topology

Assigns a weight to the Topology performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-capacity

Assigns a weight to the Virtual Server performance factor when the value of the either the load-

balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

qos-vs-score

Assigns a weight to the Virtual Server Score performance factor when the value of the either the load-balancing-mode or fallback-mode options is quality-of-service. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ttl Specifies the number of seconds that the answer, once found, is valid. Once the time-to-live (TTL) expires, the client has to request name resolution again. The valid values are 0 through 4294967295; the default value is 30.

verify-member-availability

Specifies that the system verifies the availability of the members before sending a connection to those resources. The default value is enabled.

SEE ALSO

cli admin-partitions, create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsl

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-04-28 gtm pool srv(1)

gtm prober-pool

NAME

prober-pool - Configures prober pools for the Global Traffic Manager(tm).

MODULE

gtm

SYNTAX

Modify the Global Traffic Manager prober-pool component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create prober-pool [name]
modify prober-pool [name]
options:
  app-service [[string] | none]
  description [string]
  [disabled | enabled]
  load-balancing-mode [global-availability | round-robin]
  members none
  members
    [ add | delete | modify | replace-all-with ] {
      [name] {
options:
  app-service [[string] | none]
  description [string]
  [disabled | enabled]
  order [integer]
    }...
    }
}
```

```
edit prober-pool [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

```
reset-stats prober-pool
reset-stats prober-pool [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list prober-pool
list prober-pool [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config prober-pool
show running-config prober-pool [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
```

non-default-properties

```
show prober-pool
show prober-pool [name]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

DELETE
delete prober-pool [name]
```

Note: You must remove all references to a prober-pool before you can delete the prober-pool.

DESCRIPTION

You can use the prober-pool component to configure prober pool definitions on the Global Traffic Manager. You use prober pools to control which BIG-IP servers on your network are utilized by GTM to monitor the up/down state of GTM resources. Once defined, prober pools can be set to monitor whole data centers or individual servers.

EXAMPLES

```
create prober-pool my_pool members add {
bigip-dallas
bigip-london
}
```

Creates a Global Traffic Manager prober pool with two members bigip-dallas and bigip-london. Members are selected using the global-availability load balancing method.

```
delete prober-pool my_pool
```

Deletes the prober pool my_pool.

```
show prober-pool
```

Displays statistics for all prober pools.

```
list prober-pool my_pool
```

Displays settings of prober pool my_pool.

OPTIONS

app-service

Specifies the name of the application service to which this prober pool belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this prober pool. Only the application service can modify or delete this prober pool.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

load-balancing-mode

Specifies the load balancing mode that the system uses to select members of this pool. The default value is global-availability.

The options are:

global-availability

Specifies that the Global Traffic Manager selects the first available pool member in the order in which they are listed.

round-robin

Specifies that the Global Traffic Manager selects members using a circular, sequential pattern among available pool members.

members

Specifies the BIG-IP server names of the pool members. The default value is none.

You can also use the following options with prober pool members:

app-service

Specifies the name of the application service to which this prober pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this prober pool member. Only the application service can modify or delete this prober pool member.

description

User defined description.

[enabled | disabled]

Specifies whether this pool member is available to issue probes. The default value is enabled.

order

Specifies the order number of the pool member. The system uses this number with load balancing

methods that involve prioritizing pool members by listed order.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, gtm server, gtm datacenter, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 gtm prober-pool(1)

gtm region

NAME

region - Configures a Global Traffic Manager(tm) region.

MODULE

gtm

SYNTAX

Configure the region component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

create region [name]

modify region [name]

options:

app-service [[string] | none]

description [string]

[name]

region-members

options:

app-service [[string] | none]

continent [Africa | Antarctica | Asia | Australia | Europe

| North America | South America | unknown]

country [two-letter abbreviation of country name]

datacenter [name]

geoip-isp [name]

isp [AOL | BeijingCNC | CNC | ChinaTelecom | Comcast | Earthlink

| ShanghaiCNC | ShanghaiTelecom]

not [continent | country | datacenter | isp | pool | region-name

| subnet]

pool [name]

region-name [name]

state [name]

subnet

edit region [[[name] | [glob] | [regex]] ...]

options:

all-properties

DISPLAY

list region

list region [[[name] | [glob] | [regex]] ...]

show running-config region

show running-config region [[[name] | [glob] | [regex]] ...]

options:

all-properties

one-line

DELETE

delete region [name]

DESCRIPTION

You can use the region component to create, display, modify, or delete a region. A region is a customized collection of topologies with which you can extend the topology functionality by defining specific geographical regions that have meaning for your network.

EXAMPLES

create region my_region continent Australia

Creates a region named my_region to populate with resources on the continent of Australia.

list region

Displays properties for all regions.

delete region my_region

Deletes the region named my_region.

OPTIONS

app-service

Specifies the name of the application service to which the region belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the region. Only the application service can modify or delete the region.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

region-members

Specifies the members that you want to add to, delete from, replace-all-with, or modify for this region.

You can specify the following options for region members:

app-service

Specifies the name of the application service to which the region member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the region member. Only the application service can modify or delete the region member.

continent

Specifies the name of a continent.

country

Specifies the two-letter abbreviation of a country. Use the command completion feature to view the numerous options.

datacenter

Specifies the name of an existing data center.

geoiP-isp

An ISP whose IP Address allocation range should be used in matching topologies. Any ISP string may be given as long as it matches the string in the GeoIP-ISP Database. Case is irrelevant.

isp Specifies the name of an Internet service provider.

not Specifies region-members to exclude from this region.

pool Specifies the name of an existing pool.

region-name

Specifies the name of an existing region.

state

Specifies the name of an existing state.

subnet

Specifies an existing subnet.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2015. All rights reserved.

gtm rule

NAME

rule - Opens an editor in which you can configure iRules(r) for traffic management system configuration.

MODULE

gtm

SYNTAX

Configure the rule component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create rule [name]
```

```
modify rule [name]
```

```
option:
```

```
  metadata
```

```
    [add | delete | modify] {
```

```
[metadata_name] {
```

```
  value [ "value content" ]
```

```
  persist [ true | false ]
```

```
  }
```

```
}
```

```
edit rule [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list rule
```

```
list rule [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config rule
```

```
show running-config rule [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
  all-properties
```

```
  non-default-properties
```

```
show rule
```

```
show rule [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
```

```
  field-fmt
```

DELETE

```
delete rule [name]
```

Note: You can also delete metadata associated with an iRule. See the example section for detail.

DESCRIPTION

You can use iRules to direct traffic not only to specific pools, but also to individual pool members, including port numbers and URI paths, either to implement persistence or to meet specific load balancing requirements. The syntax that you use to write iRules is based on the Tools Command Language (Tcl) programming standard. Thus, you can use many of the standard Tcl commands, plus a robust set of extensions that the BIG-IP(r) local traffic management system provides to help you further increase load balancing efficiency.

For information about standard Tcl syntax, see <http://tmml.sourceforge.net/doc/tcl/index.html>. For a list of Tcl commands that have been disabled within the traffic management system and therefore cannot be used in the traffic management system, see the Configuration Guide for BIG-IP(r) Local Traffic Management(r). This guide is available at <https://support.f5.com>.

EXAMPLES

```
edit rule my_irule
```

Opens the vi editor in which you can edit the iRule named my_irule. Note that after you close the editor, you must run the command sequence save config to save the configuration changes to the stored configuration files.

The following are example iRules for the Global Traffic Manager(tm).

```
when DNS_REQUEST {
if {[IP::addr [IP::remote_addr]/24 equals 10.10.1.0/24] }
  {cname cname.siterequest.com } else { host 10.20.20.20}}
```

Specifies that requests from 10.10.1.0/24 be directed to cname.siterequest.com, and all other requests be directed to 10.20.20.20.

```
when DNS_REQUEST {
if {[whereis [IP::remote_addr]] contains "Asia"}
  {pool asia_pool} else {pool general_pool}}
```

Specifies that requests that originate in Asia be directed to the pool named asia_pool, and that all other requests be directed to the pool named general_pool.

metadata is the user defined key/value pair

Adds new metadata to named my_meta and modifies existing metadata named my_meta2 for the iRule named my_irule.

```
modify rule my_irule {
when DNS_REQUEST {}
metadata replace-all-with {
my_meta { persist false value "hello" }
my_meta2 { persist false value "hello 2" }
} }
```

Deletes metadata named my_meta from the iRule named my_irule.

```
modify rule my_irule {
when RULE_INIT {}
definition-checksum 7c0dba9aa53e8959042c6cfe041d3d11
metadata delete { my_meta } }
```

OPTIONS

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

metadata
Specifies a user-defined key/value pair.

name Specifies a unique name for the component. This option is required.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

edit, glob, list, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-12-21 gtm rule(1)

gtm server

NAME

server - Configures servers for the Global Traffic Manager(tm).

MODULE

gtm

SYNTAX

Configure the server component within the gtm module using the syntax shown in the following sections.

```
CREATE/MODIFY
create server [name]
modify server [name]
options:
app-service [[string] | none]
datacenter [name]
description [string]
devices
[add | delete | modify | replace-all-with] {
[name] {
addresses
[add | delete | modify | replace-all-with] {
[ip address] {
options:
app-service [[string] | none]
description [string]
explicit-link-name [none | [name] ]
translation [ip address]
}
}
app-service [[string] | none]
description [string]
}
app-service [[string] | none]
description [string]
}
[disabled | enabled]
```

```

expose-route-domains [no | yes]
iq-allow-path [no | yes]
iq-allow-service-check [no | yes]
iq-allow-snmp [no | yes]
iquery-cipher-list [string]
iquery-minimum-tls-version [string]
limit-cpu-usage [integer]
limit-cpu-usage-status [disabled | enabled]
limit-mem-avail [integer]
limit-mem-avail-status [disabled | enabled]
limit-max-bps [integer]
limit-max-bps-status [disabled | enabled]
limit-max-connections [integer]
limit-max-connections-status [disabled | enabled]
limit-max-pps [integer]
limit-max-pps-status [disabled | enabled]
link-discovery [disabled | enabled]
metadata
  [add | delete | modify] {
    [metadata_name ... ] {
value [ "value content" ]
persist [true | false]
    }
  }
monitor [none | [name] [and [name]] ... ]
monitor min [integer] of { [name]... }
monitor require [integer] from [integer] { [name] }
prober-fallback [ inherit | any-available | inside-datacenter | outside-datacenter | pool | none ]
prober-pool [none | name]
prober-preference [ inherit | inside-datacenter | outside-datacenter | pool ]
product [name]
virtual-server-discovery [disabled | enabled]
virtual-servers none
virtual-servers
  [add | delete | replace-all-with] {
    [vs-name] {
options:
  app-service [[string] | none]
  depends-on none
  depends-on
    [add | delete | replace-all-with] {
      [server_name:vs-name]...
    }
  description [string]
  destination [ipv4_address:port | ipv6_address.port]
  [disabled | enabled]
  explicit-link-name [none | [name] ]
  limit-max-bps [integer]
  limit-max-bps-status [disabled | enabled]
  limit-max-connections [integer]
  limit-max-connections-status [disabled | enabled]
  limit-max-pps [integer]
  limit-max-pps-status [disabled | enabled]
  ltm-name [name]
  monitor [none | [name] [and [name]] ... ]
  monitor min [integer] of { [name]... }
  monitor require [integer] from [integer] { [name] }
  translation-address [ip address]
  translation-port [ [integer] | [name] ]
  }
}

```

```
edit server [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
all-properties
non-default-properties
one-line
```

```
reset-stats server
```

```
reset-stats server [ [ [name] | [glob] | [regex] ] ... ]
```

```
DISPLAY
```

```
list server
```

```
list server [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config server
```

```
show running-config server [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
all-properties
non-default-properties
partition
```

```
show server
```

```
show server [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt
```

DELETE
delete server [name]

DESCRIPTION

You can use the server component to configure a Global Traffic Manager server.

EXAMPLES

```
create server my_server devices add { my_device { addresses add { 10.10.1.1 } } } datacenter my_datacenter
```

Creates a server named my_server in my_datacenter with a self IP address of 10.10.1.1.

```
modify server my_server virtual-servers add {myVs { address 10.10.10.2:80 } }
```

Adds the virtual server myVs with an IP address of 10.10.10.2:80 as a resource to the server named my_server.

```
list server non-default-properties
```

Displays all non-default properties for all servers.

```
delete server my_server
```

Deletes the server named my_server.

```
show server my_server detail
```

Shows the link associated with each server IP address for my_server.

```
show server all detail
```

Shows the link assignments for all servers in the system.

```
show server my_server virtual-servers
```

Shows the regular server information as well as any virtuals on my_server. The link associated with a virtual server is displayed, or --- is shown to indicate that the virtual is not using a configured link.

OPTIONS

app-service

Specifies the name of the application service to which the server belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the server. Only the application service can modify or delete the server.

datacenter

Specifies the data center to which the server belongs. This option is required for the command create.

description

User defined description.

detail

The detail option is used with the show display command. This shows device statistics, virtual server statistics, and links associated with the server. Only the device addresses that have an associated link are displayed. If this server is not using any links, no link assignment information is printed. By default, links are automatically matched to device addresses according to the smallest subnet match. Explicit links may also be defined. How this link was assigned is displayed in the Link Assignment column: auto means that the system automatically assigned this link, and explicit means that the link was explicitly set by the user.

devices

Specifies the names of the devices that represent this server. Every device must have at least one address. The options are:

addresses

Specifies the IP addresses that represent the device. If GTM configuration synchronization is enabled and all existing addresses for a device are being replaced, new addresses should be added and synchronized before old addresses are removed, otherwise the changes may fail to synchronize. Alternatively, the address configuration changes can be performed on each GTM. The options are:

app-service

Specifies the name of the application service to which the address belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the address. Only the application service can modify or delete the address.

description

User defined description.

explicit-link-name

Specifies the explicit link name for the device. The default value is none.

translation

Specifies the internal IP address that corresponds to the external IP address of this device. The default value is ::.

app-service

Specifies the name of the application service to which the device belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the device. Only the application service can modify or delete the

device.

description
User defined description.

[disabled | enabled]
Enables or disables the server. The default value is enabled.

expose-route-domains
Allow the GTM server to auto-discover LTM virtual servers from all route domains. The default value is no.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

iq-allow-path
Specifies whether the Global Traffic Manager uses this BIG-IP(r) system to conduct a path probe before delegating traffic to it. The default value is yes.

iq-allow-service-check
Specifies whether the Global Traffic Manager uses this BIG-IP system to conduct a service check probe before delegating traffic to it. The default value is yes.

iq-allow-snmp
Specifies whether the Global Traffic Manager uses this BIG-IP system to conduct an SNMP probe before delegating traffic to it. The default value is yes.

iquery-cipher-list
This is a ":" separated list of cipher specifications as accepted by the "openssl ciphers" command. OpenSSL will use the cipher list to negotiate a mutually acceptable cipher with the server during iQuery connection setup. Setting this value on the server will override the value inherited from the global settings.

iquery-minimum-tls-version
This is a string to specify the minimum TLS version that will be offered by the client (GTM) during iQuery connection negotiation. Setting this value on the server will override the value inherited from the global settings.

limit-cpu-usage
For a server configured as a generic host, specifies the percent of CPU usage, otherwise has no effect. If percent of CPU usage goes above the limit, the system marks the server as unavailable.

limit-cpu-usage-status
Enables or disables the limit-cpu-usage option for this server. Only has an effect on a server configured as a generic host. The default value is disabled.

limit-mem-avail
For a server configured as a generic host, specifies the available memory required by the virtual servers on the server. If available memory falls below this limit, the system marks the server as unavailable.

limit-mem-avail-status
Enables or disables the limit-mem-avail option for this server. Only used on a server configured as a generic host. The default value is disabled.

limit-max-bps
Specifies the maximum allowable data throughput rate, in bits per second, for this server. If the network traffic volume exceeds this limit, the system marks the server as unavailable.

limit-max-bps-status
Enables or disables the limit-max-bps option for this server. The default value is disabled.

limit-max-connections
Specifies the maximum number of concurrent connections, combined, for this server. If the connections exceed this limit, the system marks the server as unavailable.

limit-max-connections-status
Enables or disables the limit-max-connections option for this server. The default value is disabled.

limit-max-pps
Specifies the maximum allowable data transfer rate, in packets per second, for this server. If the network traffic volume exceeds this limit, the system marks the server as unavailable.

limit-max-pps-status
Enables or disables the limit-max-pps option for this server. The default value is disabled.

link-discovery
Specifies whether the system auto-discovers the links for this server. The default value is disabled. The options are:

disabled
Specifies that the system does not auto-discover the links that are available for the server.

enabled
Specifies that the system auto-discovers the links that are configured on the server. With this option, the system automatically adds, deletes, and modifies link settings in the configuration.

enabled-no-delete

Specifies that the system auto-discovers the links that are configured on the server. With this option, the system automatically adds and modifies link settings in the configuration, but does not delete them. This option is useful when you regularly take links in and out of service.

metadata

Specifies user-defined data to associate with a server. By default the persist attribute is set to true. This means the data is saved into the configuration file.

monitor

Specifies the health monitors that the system uses to determine whether this server is available for load balancing. Multiple monitors may be specified with the and keyword. The min keyword is used to specify the minimum number of monitors that must succeed for this server to be declared up. The require keyword is used to specify the minimum number of probes that must succeed for this server to be declared up and the number of probes that should be used.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the object resides.

prober-fallback

Specifies the type of prober to use to monitor this server's resources when the preferred type is not available. If this value is specified, it overrides the prober fallback value on this server's data center. The default value is inherit.

prober-pool

Specifies the name of a prober pool to use to monitor this server's resources when either the prober-preference or prober-fallback value is pool. If neither the prober-preference or prober-fallback value is pool, the prober-preference and prober-fallback values are set to pool and any-available. If this value is specified, it overrides any prober pool set on this server's data center. The default value is none.

prober-preference

Specifies the type of prober to use to monitor this server's resources. If this value is specified, it overrides the prober preference value on this server's data center. The default value is inherit.

product

Specifies the server type. The server type determines the metrics that the system can collect from the server. Use the command completion feature to view the types of servers that are available.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

virtual-server-discovery

Specifies whether the system auto-discovers the virtual servers for this server. The default value is disabled. The options are:

disabled

Specifies that the system does not auto-discover the virtual servers that are configured on the server. With this option, you must configure the virtual servers for this server.

enabled

Specifies that the system auto-discovers the virtual servers that are configured on the server. With this option, the system automatically adds, deletes, and modifies virtual server settings in the configuration.

enabled-no-delete

Specifies that the system auto-discovers the virtual servers that are configured on the server. With this option, the system automatically adds and modifies virtual server settings in the configuration, but does not delete them. This option is useful when you regularly take virtual servers in and out of service.

virtual-servers

Specifies the name of the virtual servers that are resources for this server. You can include the following options for virtual servers.

app-service

Specifies the name of the application service to which the virtual server belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the virtual server. Only the application service can modify or delete the virtual server.

depends-on

Specifies the vs-name of the server on which this virtual server depends.

description

User defined description.

destination

Specifies the IP address and port of the virtual server.

[disabled | enabled]

Specifies whether this virtual server is available for load balancing. The default value is enabled.

explicit-link-name

Specifies the explicit link name for the virtual server. The default value is none.

limit-max-bps

Specifies the maximum allowable data throughput rate, in bits per second, for this virtual server. If the network traffic volume exceeds this value, the system marks the virtual server as unavailable. The default value is 0 (zero).

limit-max-bps-status

Enables or disables the limit-max-bps option for this virtual server. The default value is disabled.

limit-max-connections

Specifies the number of current connections allowed for this virtual server. If the current connections exceed this value, the system marks this virtual server as unavailable. The default value is 0 (zero).

limit-max-connections-status

Enables or disables the limit-max-connections option for this virtual server. The default value is disabled.

limit-max-pps

Specifies the maximum allowable data transfer rate, in packets per second, for this virtual server. If the network traffic volume exceeds this limit, the system marks the virtual server as unavailable. The default value is 0 (zero).

limit-max-pps-status

Enables or disables the limit-max-pps option for this virtual server. The default value is disabled.

ltn-name

The virtual server name found on the LTM. Useful for differentiating between virtuals with same IP and port, but different protocols. The ltn-name used in probe requests.

monitor

Specifies the monitor you want to assign to this virtual server. Multiple monitors may be specified with the and keyword. The min keyword is used to specify the minimum number of monitors that must succeed for this server to be declared up. The require keyword is used to specify the minimum number of probes that must succeed for this server to be declared up and the number of probes that should be used. The default value is none.

translation-address

Specifies the public address that this virtual server translates into when the Global Traffic Manager communicates between the network and the Internet. The default value is ::.

translation-port

Specifies the translation port number or service name for the virtual server, if necessary. The default value is 0.

SEE ALSO

create, delete, edit, glob, gtm datacenter, gtm link, gtm prober-pool, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2018-03-09 gtm server(1)

gtm topology

NAME

topology - Configures a topology statement.

MODULE

gtm

SYNTAX

Configure the topology component within the gtm module using the syntax shown in the following sections.

CREATE

create topology

options:

app-service [[string] | none]

description [[string] | none]

ldns: [continent | country | geoip-isp | isp | not | region | state | subnet]

order [integer]

score [integer]

server: [continent | country | datacenter | geoip-isp | isp | not | pool | region | state | subnet]

MODIFY

modify topology [[name] | [glob] | [regex]] ...]

options:

app-service [[string] | none]
description [[string] | none]
order [integer]
score [integer]

EDIT

edit topology [[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list topology

list topology [[name] | [glob] | [regex]] ...]

show running-config topology

show running-config topology [[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line

DELETE

delete topology all

delete topology

[ldns: [identifier] [value] server: [identifier] [value]]

DESCRIPTION

You can use the topology component to configure a topology statement. A topology statement is a set of characteristics that identify the origin of a given name resolution request.

EXAMPLES

```
create topology ldns: country US server: datacenter DC1 score 30
```

Creates a topology statement that specifies that the Global Traffic Manager routes any traffic coming from the United States to the datacenter named DC1. Note that the weight of this topology item for load balancing is 30.

```
delete topology ldns: country US server: datacenter DC1
```

Deletes the topology statement mentioned in the previous example.

OPTIONS

app-service

Specifies the name of the application service to which the topology belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the topology. Only the application service can modify or delete the topology.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ldns:

Specifies the criteria that the Global Traffic Manager uses when matching requests from LDNS servers.

continent

A continent whose IP address allocation range should be used in matching topologies

country

A country whose IP address allocation range should be used in matching topologies

datacenter

A data center to be used in matching topologies

geoip-isp

An ISP whose IP Address allocation range should be used in matching topologies. Any ISP string may be given as long as it matches the string in the GeoIP-ISP Database. Case is irrelevant.

isp An ISP whose IP address allocation range should be used in matching topologies

not Specify a region member to exclude from the region

pool A pool to be used in matching topologies

region

Another region to be used in matching topologies

state

A state whose IP address allocation range should be used in matching topologies

subnet

A subnet to be used in matching topologies

order

Specify the order number of this topology item. Should not be specified if the Longest Match option is selected. Must be a value between 1 and the number of topology items inclusive (including any new topology items). If an existing topology item specifies the same order number, the order numbers for that item and all subsequent items are incremented.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

score

Specifies the weight of the topology item.

server:

Specifies the server to which the Global Traffic Manager routes requests.

continent

A continent whose IP address allocation range should be used as an LDNS routing destination

country

A country whose IP address allocation range should be used as an LDNS routing destination

datacenter

A data center to be used as an LDNS routing destination

geop-isp

An ISP whose IP Address allocation range should be used in matching topologies. Any ISP string may be given as long as it matches the string in the GeoIP-ISP Database. Case is irrelevant.

isp An ISP whose IP address allocation range should be used as an LDNS routing destination

not Specify an item to exclude from the group

pool A pool to be used as an LDNS routing destination

region

Another region to be used as an LDNS routing destination

state

A state whose IP address allocation range be used as an LDNS routing destination

subnet

A subnet to be used as an LDNS routing destination

SEE ALSO

create, delete, edit, glob, gtm server, list, regex, show, tmsh,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2015. All rights reserved.

BIG-IP 2018-06-14 gtm topology(1)

gtm traffic

NAME

traffic - Displays traffic statistics for the Global Traffic Manager(tm).

MODULE

gtm

SYNTAX

Configure the traffic component within the gtm module using the syntax in the following section.

DISPLAY

show traffic

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DESCRIPTION

You can use the traffic component to display traffic statistics, including those for IPv4 and IPv6 requests, current Local Domain Name System (LDNS) servers, and current paths.

SEE ALSO

show, sys tmm-traffic, tmmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 gtm traffic(1)

gtm wideip a

NAME

a - Configures a wide IP that accepts A queries.

MODULE

gtm wideip

SYNTAX

Configure the wideip a component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create wideip a [name]
modify wideip a [name]
options:
  aliases [name...name]
  app-service [[string] | none]
  description [string]
  [disabled | enabled]
  failure-rcode [formerr | noerror | notimpl | nxdomain | refused | servfail]
  failure-rcode-ttl [integer]
  failure-rcode-response [disabled | enabled]
  last-resort-pool [type name]
  load-balancing-decision-log-verbosity [[pool-member-selection | pool-member-traversal | pool-selection | pool-traversal] | none]
  minimal-response [disabled | enabled]
  metadata none
  metadata
    [add | delete | modify | replace-all-with] {
[metadata_name ... ] {
  persist [ true | false ]
  value [ "value content" ]
}
}
  persistence [disabled | enabled]
  persist-cidr-ipv4 [integer]
  persist-cidr-ipv6 [integer]
  pool-lb-mode [global-availability | ratio | round-robin | topology]
  pools none
  pools
    [add | delete | modify | replace-all-with] {
[pool name]...
}
  pools-cname none
  pools-cname
    [add | delete | modify | replace-all-with] {
[pool name]...
}
  rules none
  rules {
[rule name]
...
}
  topology-prefer-edns0-client-subnet [disabled | enabled]
  ttl-persistence [integer]
```

```
edit wideip a [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
reset-stats wideip a
reset-stats wideip a [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list wideip a
list wideip a [ [name] | [glob] | [regex] ] ... ]
show running-config wideip a
```

show running-config wideip a [[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

show wideip a

show wideip a [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
(detail | global)
field-fmt

DELETE

delete wideip a [all | [name]]

DESCRIPTION

You can use the wideip component to create, modify, display, or delete a wide IP that responds to A queries. An A wide IP is a mapping of a fully-qualified domain name (FQDN) to a set of IPv4 virtual servers that host the domain's content, such as a web site or an e-commerce site.

EXAMPLES

```
create wideip a www.my_wide_ip.com
```

Creates a wide IP named www.my_wide_ip.com.

```
delete wideip a www.my_wide_ip.com
```

Deletes the wide IP named www.my_wide_ip.

OPTIONS

aliases

Specifies alternate domain names for the web site content you are load balancing. You can use two different wildcard characters, asterisk (*) and question mark (?), to represent one or more characters. The default value is none.

app-service

Specifies the name of the application service to which this wide ip belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this wide ip. Only the application service can modify or delete this wide ip.

description

User defined description. Note: Advanced search on the GUI's Wideip list page can be turned on/off by modifying the DB variable ui.advancedsearch via the tmsh command "modify sys db ui.advancedsearch value true/false". This will result in a new description column and the inclusion of that field in the search.

[disabled | enabled]

Specifies whether the wide IP and its resources are available for load balancing.

failure-rcode

Specifies the DNS RCODE used when failure-rcode-response is enabled. Default is noerror. Options include noerror (no type exists at this name), formerr (format error in query), servfail (unable to process query), nxdomain (name does not exist), notimpl (no support for this kind of query), and refused (refuse to process based on policy). If failure-rcode-ttl is non-zero, only the Authority section of the noerror or nxdomain response will include a SOA record.

failure-rcode-response

When enabled, specifies that the system returns a RCODE response to Wide IP requests after exhausting all load-balancing methods. This response is an authoritative empty answer from the system to A record requests. With this option enabled, the system responds faster to A requests for which it does not have A records configured. The default value is disabled.

failure-rcode-ttl

Specifies the negative caching TTL of the SOA for the RCODE response. The default is 0, meaning no SOA is included (i.e. no caching).

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

last-resort-pool

Specifies a last resort pool to use when load balancing requests for this wide IP. Any A or CNAME pool type is allowed. The default value is none.

load-balancing-decision-log-verbosity

Specifies the amount of detail logged when making load balancing decisions. This is used for debugging purpose only. Performance will be affected if the value is not none. Please reset it back to none after done debugging. With the option pool-selection, the log will contain pool load balancing algorithm details. This includes common actions taken to a set of pools (for example, whether all pools reset the ratio counter during the algorithm) and the result of the load balancing algorithm (for example, whether a pool is finally selected and the reason if applicable). With the option pool-traversal, the log will contain details of all pools traversed during load balancing. With the option pool-member-selection, the log will contain pool member load balancing algorithm details. This includes common actions taken to a set of pool members and the result of the load balancing algorithm. With the option pool-member-traversal, the log will contain details of all pool members traversed during load balancing. The default value is none.

minimal-response

Specifies GTM will form the smallest allowable DNS response to a query. Typically, this will be a single resource record in the answer section. When set to disabled, GTM will attempt to chase CNAME chains, if required, to obtain the ultimate answer, and it will attempt to add address resource records to the additional section of the response for each answer when needed. The default value is enabled.

metadata

Specifies user-defined data to associate with a server. By default the persist attribute is set to true. This means the data is saved into the configuration file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

persistence

When enabled, specifies that when a local DNS server makes repetitive requests on behalf of a client, the system reconnects the client to the same resource as previous requests. The default value is disabled.

persist-cidr-ipv4

Specifies a mask used to group IPv4 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

persist-cidr-ipv6

Specifies a mask used to group IPv6 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

pools

Configures the pools the system uses when load balancing requests for this wide IP. The default value is none.

pools-cname

Configures the CNAME pools the system uses when load balancing requests for this wide IP. The default value is none.

pool-lb-mode

Specifies the load balancing method used to select a pool in this wide IP. This option is relevant only when multiple pools are configured for this wide IP. The default value is round-robin.

The available load balancing methods are:

global-availability

Specifies that the system selects a pool by following the order of the Pool list. The system repeatedly selects the first pool in the list for as long as its status is available. If the pool becomes unavailable for any reason, the system then repeatedly selects the next pool in the list, and so on.

ratio

Specifies that the system selects a pool based on the ratio that you assign to the pool.

round-robin

Specifies that the system selects pools sequentially.

topology

Specifies that the system selects a pool based on topology information in the incoming LDNS request. Note that this load balancing method works only if you have configured a topology statement.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rules

Specifies the iRules(r) that this wide IP uses for load balancing decisions. 'when' clauses for each event are grouped across all iRules(r) on this wide IP. For each event, clauses are evaluated in the listed rules order. The default value is none.

ttd-persistence

Specifies, in seconds, the length of time for which a persistence entry is valid. This value can range from 0 through 4294967295 seconds. The default value is 3600.

topology-prefer-edns0-client-subnet

Specifies, when set to enabled, that this wide IP should use the edns0 client subnet option (if one exists) instead of the source address when using topology load balancing. When this option is set to disabled or if the query did not contain a client subnet option, the system will fall back to the source address.

This setting has no effect when the global setting, configured under gtm global-settings load-balancing topology-prefer-edns0-client-subnet [disabled | enabled], is set to enabled. When either setting is enabled then this feature will be enabled.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, reset-stats, show, tmsd

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

gtm wideip aaaa

NAME

aaaa - Configures a wide IP that accepts AAAA queries.

MODULE

gtm wideip

SYNTAX

Configure the wideip aaaa component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create wideip aaaa [name]
modify wideip aaaa [name]
options:
  aliases [name...name]
  app-service [[string] | none]
  description [string]
  [disabled | enabled]
  failure-rcode [formerr | noerror | notimpl | nxdomain | refused | servfail]
  failure-rcode-ttl [integer]
  failure-rcode-response [disabled | enabled]
  last-resort-pool [type name]
  load-balancing-decision-log-verbosity [[pool-member-selection | pool-member-traversal | pool-selection | pool-traversal] | none]
  minimal-response [disabled | enabled]
  metadata none
  metadata
    [add | delete | modify | replace-all-with] {
[metadata_name ... ] {
  persist [ true | false ]
  value [ "value content" ]
}
}
  persistence [disabled | enabled]
  persist-cidr-ipv4 [integer]
  persist-cidr-ipv6 [integer]
  pool-lb-mode [global-availability | ratio | round-robin | topology]
  pools none
  pools
    [add | delete | modify | replace-all-with] {
[pool name]...
}
  pools-cname none
  pools-cname
    [add | delete | modify | replace-all-with] {
[pool name]...
}
  rules none
  rules {
[rule name]
...
}
  topology-prefer-edns0-client-subnet [disabled | enabled]
  ttl-persistence [integer]
```

```
edit wideip aaaa [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

```
reset-stats wideip aaaa
reset-stats wideip aaaa [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list wideip aaaa
list wideip aaaa [ [name] | [glob] | [regex] ] ... ]
show running-config wideip aaaa
show running-config wideip aaaa [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
show wideip aaaa
```

show wideip aaaa [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
(detail | global)
field-fmt

DELETE
delete wideip aaaa [all | [name]]

DESCRIPTION

You can use the wideip component to create, modify, display, or delete a wide IP that responds to AAAA queries. An AAAA wide IP is a mapping of a fully-qualified domain name (FQDN) to a set of IPv6 virtual servers that host the domain's content, such as a web site or an e-commerce site.

EXAMPLES

```
create wideip aaaa www.my_wide_ip.com
```

Creates a AAAA wide IP named www.my_wide_ip.com.

```
delete wideip aaaa www.my_wide_ip.com
```

Deletes the AAAA wide IP named www.my_wide_ip.

OPTIONS

aliases

Specifies alternate domain names for the web site content you are load balancing. You can use two different wildcard characters, asterisk (*) and question mark (?), to represent one or more characters. The default value is none.

app-service

Specifies the name of the application service to which this wide ip belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this wide ip. Only the application service can modify or delete this wide ip.

description

User defined description. Note: Advanced search on the GUI's Wideip list page can be turned on/off by modifying the DB variable ui.advancedsearch via the tmsb command "modify sys db ui.advancedsearch value true/false". This will result in a new description column and the inclusion of that field in the search.

[disabled | enabled]

Specifies whether the wide IP and its resources are available for load balancing.

failure-rcode

Specifies the DNS RCODE used when failure-rcode-response is enabled. Default is noerror. Options include noerror (no type exists at this name), formerr (format error in query), servfail (unable to process query), nxdomain (name does not exist), notimpl (no support for this kind of query), and refused (refuse to process based on policy). If failure-rcode-ttl is non-zero, only the Authority section of the noerror or nxdomain response will include a SOA record.

failure-rcode-response

When enabled, specifies that the system returns a RCODE response to Wide IP requests after exhausting all load-balancing methods. This response is an authoritative empty answer from the system to AAAA record requests. With this option enabled, the system responds faster to AAAA requests for which it does not have AAAA records configured. The default value is disabled.

failure-rcode-ttl

Specifies the negative caching TTL of the SOA for the RCODE response. The default is 0, meaning no SOA is included (i.e. no caching).

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

last-resort-pool Specifies a last resort pool to use when load balancing requests for this wide IP. Any AAAA or CNAME pool type is allowed. The default value is none.

load-balancing-decision-log-verbosity

Specifies the amount of detail logged when making load balancing decisions. This is used for debugging purpose only. Performance will be affected if the value is not none. Please reset it back to none after done debugging. With the option pool-selection, the log will contain pool load balancing algorithm details. This includes common actions taken to a set of pools (for example, whether all pools reset the ratio counter during the algorithm) and the result of the load balancing algorithm (for example, whether a pool is finally selected and the reason if applicable). With the option pool-traversal, the log will contain details of all pools traversed during load balancing. With the option pool-member-selection, the log will contain pool member load balancing algorithm details. This includes common actions taken to a set of pool members and the result of the load balancing algorithm. With the option pool-member-traversal, the log will contain details of all pool members traversed during load balancing. The default value is none.

minimal-response

Specifies GTM will form the smallest allowable DNS response to a query. Typically, this will be a single resource record in the answer section. When set to disabled, GTM will attempt to chase CNAME chains, if required, to obtain the ultimate answer, and it will attempt to add address resource records to the additional section of the response for each answer when needed. The default value is enabled.

metadata

Specifies user-defined data to associate with a server. By default the persist attribute is set to true. This means the data is saved into the configuration file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

persistence

When enabled, specifies that when a local DNS server makes repetitive requests on behalf of a client, the system reconnects the client to the same resource as previous requests. The default value is disabled.

persist-cidr-ipv4

Specifies a mask used to group IPv4 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

persist-cidr-ipv6

Specifies a mask used to group IPv6 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

pools

Configures the pools the system uses when load balancing requests for this wide IP. The default value is none.

pools-cname

Configures the CNAME pools the system uses when load balancing requests for this wide IP. The default value is none.

pool-lb-mode

Specifies the load balancing method used to select a pool in this wide IP. This option is relevant only when multiple pools are configured for this wide IP. The default value is round-robin.

The available load balancing methods are:

global-availability

Specifies that the system selects a pool by following the order of the Pool list. The system repeatedly selects the first pool in the list for as long as its status is available. If the pool becomes unavailable for any reason, the system then repeatedly selects the next pool in the list, and so on.

ratio

Specifies that the system selects a pool based on the ratio that you assign to the pool.

round-robin

Specifies that the system selects pools sequentially.

topology

Specifies that the system selects a pool based on topology information in the incoming LDNS request. Note that this load balancing method works only if you have configured a topology statement.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rules

Specifies the iRules(r) that this wide IP uses for load balancing decisions. 'when' clauses for each event are grouped across all iRules(r) on this wide IP. For each event, clauses are evaluated in the listed rules order. The default value is none.

ttl-persistence

Specifies, in seconds, the length of time for which a persistence entry is valid. This value can range from 0 through 4294967295 seconds. The default value is 3600.

topology-prefer-edns0-client-subnet

Specifies, when set to enabled, that this wide IP should use the edns0 client subnet option (if one exists) instead of the source address when using topology load balancing. When this option is set to disabled or if the query did not contain a client subnet option, the system will fall back to the source address.

This setting has no effect when the global setting, configured under gtm global-settings load-balancing topology-prefer-edns0-client-subnet [disabled | enabled], is set to enabled. When either setting is enabled then this feature will be enabled.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

gtm wideip cname

NAME

cname - Configures a wide IP that accepts CNAME queries.

MODULE

gtm wideip

SYNTAX

Configure the wideip cname component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

create wideip cname [name]

modify wideip cname [name]

options:

aliases [name...name]

app-service [[string] | none]

description [string]

[disabled | enabled]

failure-rcode [formerr | noerror | notimpl | nxdomain | refused | servfail]

failure-rcode-ttl [integer]

failure-rcode-response [disabled | enabled]

last-resort-pool [type name]

load-balancing-decision-log-verbosity [[pool-member-selection | pool-member-traversal | pool-selection | pool-traversal] | none]

minimal-response [disabled | enabled]

metadata none

metadata

[add | delete | modify | replace-all-with] {

[metadata_name ...] {

persist [true | false]

value ["value content"]

}

}

persistence [disabled | enabled]

persist-cidr-ipv4 [integer]

persist-cidr-ipv6 [integer]

pool-lb-mode [global-availability | ratio | round-robin | topology]

pools none

pools

[add | delete | modify | replace-all-with] {

[pool name]...

}

rules none

rules {

[rule name]

...

}

topology-prefer-edns0-client-subnet [disabled | enabled]

ttl-persistence [integer]

edit wideip cname [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats wideip cname

reset-stats wideip cname [[name] | [glob] | [regex]] ...]

DISPLAY

list wideip cname

list wideip cname [[name] | [glob] | [regex]] ...]

show running-config wideip cname

show running-config wideip cname [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show wideip cname

show wideip cname [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

(detail | global)

field-fmt

DELETE

delete wideip cname [all | [name]]

DESCRIPTION

You can use the wideip component to create, modify, display, or delete a wide IP that responds to CNAME queries. A CNAME wide IP is a mapping of a fully-qualified domain name (FQDN) to its canonical name.

EXAMPLES

create wideip cname www.my_wide_ip.com

Creates a CNAME wide IP named www.my_wide_ip.com.

delete wideip cname www.my_wide_ip.com

Deletes the CNAME wide IP named www.my_wide_ip.

OPTIONS

aliases

Specifies alternate domain names for the web site content you are load balancing. You can use two different wildcard characters, asterisk (*) and question mark (?), to represent one or more characters. The default value is none.

app-service

Specifies the name of the application service to which this wide ip belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this wide ip. Only the application service can modify or delete this wide ip.

description

User defined description. Note: Advanced search on the GUI's Wideip list page can be turned on/off by modifying the DB variable ui.advancedsearch via the tmsh command "modify sys db ui.advancedsearch value true/false". This will result in a new description column and the inclusion of that field in the search.

[disabled | enabled]

Specifies whether the wide IP and its resources are available for load balancing.

failure-rcode

Specifies the DNS RCODE used when failure-rcode-response is enabled. Default is noerror. Options include noerror (no type exists at this name), formerr (format error in query), servfail (unable to process query), nxdomain (name does not exist), notimpl (no support for this kind of query), and refused (refuse to process based on policy). If failure-rcode-ttl is non-zero, only the Authority section of the noerror or nxdomain response will include a SOA record.

failure-rcode-response

When enabled, specifies that the system returns a RCODE response to Wide IP requests after exhausting all load-balancing methods. This response is an authoritative empty answer from the system to CNAME record requests. With this option enabled, the system responds faster to CNAME requests for which it does not have CNAME records configured. The default value is disabled.

failure-rcode-ttl

Specifies the negative caching TTL of the SOA for the RCODE response. The default is 0, meaning no SOA is included (i.e. no caching).

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

last-resort-pool

Specifies a last resort pool to use when load balancing requests for this wide IP. Any CNAME pool type is allowed. The default value is none.

load-balancing-decision-log-verbosity

Specifies the amount of detail logged when making load balancing decisions. This is used for debugging purpose only. Performance will be affected if the value is not none. Please reset it back to none after done debugging. With the option pool-selection, the log will contain pool load balancing algorithm details. This includes common actions taken to a set of pools (for example, whether all pools reset the ratio counter during the algorithm) and the result of the load balancing algorithm (for example, whether a pool is finally selected and the reason if applicable). With the option pool-traversal, the log will contain details of all pools traversed during load balancing. With the option pool-member-selection, the log will contain pool member load balancing algorithm details. This includes common actions taken to a set of pool members and the result of the load balancing algorithm. With the option pool-member-traversal, the log will contain details of all pool members traversed during load balancing. The default value is none.

minimal-response

Specifies GTM will form the smallest allowable DNS response to a query. Typically, this will be a single resource record in the answer section. When set to disabled, GTM will attempt to chase CNAME chains, if required, to obtain the ultimate answer, and it will attempt to add address resource records to the additional section of the response for each answer when needed. The default value is enabled.

metadata

Specifies user-defined data to associate with a server. By default the persist attribute is set to true. This means the data is saved into the configuration file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

persistence

When enabled, specifies that when a local DNS server makes repetitive requests on behalf of a client, the system reconnects the client to the same resource as previous requests. The default value is disabled.

persist-cidr-ipv4

Specifies a mask used to group IPv4 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

`persist-cidr-ipv6`

Specifies a mask used to group IPv6 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

`pools`

Configures the pools the system uses when load balancing requests for this wide IP. The default value is none.

`pool-lb-mode`

Specifies the load balancing method used to select a pool in this wide IP. This option is relevant only when multiple pools are configured for this wide IP. The default value is round-robin.

The available load balancing methods are:

`global-availability`

Specifies that the system selects a pool by following the order of the Pool list. The system repeatedly selects the first pool in the list for as long as its status is available. If the pool becomes unavailable for any reason, the system then repeatedly selects the next pool in the list, and so on.

`ratio`

Specifies that the system selects a pool based on the ratio that you assign to the pool.

`round-robin`

Specifies that the system selects pools sequentially.

`topology`

Specifies that the system selects a pool based on topology information in the incoming LDNS request. Note that this load balancing method works only if you have configured a topology statement.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`rules`

Specifies the `iRules(r)` that this wide IP uses for load balancing decisions. 'when' clauses for each event are grouped across all `iRules(r)` on this wide IP. For each event, clauses are evaluated in the listed rules order. The default value is none.

`ttl-persistence`

Specifies, in seconds, the length of time for which a persistence entry is valid. This value can range from 0 through 4294967295 seconds. The default value is 3600.

`topology-prefer-edns0-client-subnet`

Specifies, when set to enabled, that this wide IP should use the edns0 client subnet option (if one exists) instead of the source address when using topology load balancing. When this option is set to disabled or if the query did not contain a client subnet option, the system will fall back to the source address.

This setting has no effect when the global setting, configured under `gtm global-settings load-balancing topology-prefer-edns0-client-subnet [disabled | enabled]`, is set to enabled. When either setting is enabled then this feature will be enabled.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `gtm pool`, `list`, `modify`, `regex`, `reset-stats`, `show`, `tms`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-09-28 gtm wideip cname(1)

gtm wideip mx

NAME

`mx` - Configures a wide IP that accepts MX queries.

MODULE

`gtm wideip`

SYNTAX

Configure the `wideip mx` component within the `gtm` module using the syntax shown in the following sections.

CREATE/MODIFY

`create wideip mx [name]`

```

modify wideip mx [name]
options:
  aliases [name...name]
  app-service [[string] | none]
  description [string]
  [disabled | enabled]
  failure-rcode [formerr | noerror | notimpl | nxdomain| refused | servfail]
  failure-rcode-ttl [integer]
  failure-rcode-response [disabled | enabled]
  last-resort-pool [type name]
  load-balancing-decision-log-verbosity [[pool-member-selection | pool-member-traversal | pool-selection | pool-traversal] | none]
  minimal-response [disabled | enabled]
  metadata none
  metadata
    [add | delete | modify | replace-all-with] {
[metadata_name ... ] {
  persist [ true | false ]
  value [ "value content" ]
}
}
persistence [disabled | enabled]
  persist-cidr-ipv4 [integer]
  persist-cidr-ipv6 [integer]
  pool-lb-mode [global-availability | ratio | round-robin | topology]
  pools none
  pools
    [add | delete | modify | replace-all-with] {
[pool name]...
}
  pools-cname none
  pools-cname
    [add | delete | modify | replace-all-with] {
[pool name]...
}
  rules none
  rules {
[rule name]
...
}
topology-prefer-edns0-client-subnet [disabled | enabled]
ttl-persistence [integer]

```

```

edit wideip mx [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

```

reset-stats wideip mx
reset-stats wideip mx [ [name] | [glob] | [regex] ] ... ]

```

```

DISPLAY
list wideip mx
list wideip mx [ [name] | [glob] | [regex] ] ... ]
show running-config wideip mx
show running-config wideip mx [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
show wideip mx
show wideip mx [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  (detail | global)
  field-fmt

```

```

DELETE
delete wideip mx [all | [name] ]

```

DESCRIPTION

You can use the wideip component to create, modify, display, or delete a wide IP that responds to MX queries. An MX wide IP maps a domain's mail address to a set of hosts acting as mail servers.

EXAMPLES

```
create wideip mx mydomain.com
```

Creates an MX wide IP named mydomain.com.

```
delete wideip mx mydomain.com
```

Deletes the MX wide IP named mydomain.com.

OPTIONS

aliases

Specifies alternate domain names for the web site content you are load balancing. You can use two different wildcard characters, asterisk (*) and question mark (?), to represent one or more characters.

The default value is none.

app-service

Specifies the name of the application service to which this wide ip belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this wide ip. Only the application service can modify or delete this wide ip.

description

User defined description. Note: Advanced search on the GUI's Wideip list page can be turned on/off by modifying the DB variable ui.advancedsearch via the tmsh command "modify sys db ui.advancedsearch value true/false". This will result in a new description column and the inclusion of that field in the search.

[disabled | enabled]

Specifies whether the wide IP and its resources are available for load balancing.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

failure-rcode

Specifies the DNS RCODE used when failure-rcode-response is enabled. Default is noerror. Options include noerror (no type exists at this name), formerr (format error in query), servfail (unable to process query), nxdomain (name does not exist), notimpl (no support for this kind of query), and refused (refuse to process based on policy). If failure-rcode-ttl is non-zero, only the Authority section of the noerror or nxdomain response will include a SOA record.

failure-rcode-response

When enabled, specifies that the system returns a RCODE response to Wide IP requests after exhausting all load-balancing methods. This response is an authoritative empty answer from the system to MX record requests. With this option enabled, the system responds faster to MX requests for which it does not have MX records configured. The default value is disabled.

failure-rcode-ttl

Specifies the negative caching TTL of the SOA for the RCODE response. The default is 0, meaning no SOA is included (i.e. no caching).

last-resort-pool

Specifies a last resort pool to use when load balancing requests for this wide IP. Any MX or CNAME pool type is allowed. The default value is none.

load-balancing-decision-log-verbosity

Specifies the amount of detail logged when making load balancing decisions. This is used for debugging purpose only. Performance will be affected if the value is not none. Please reset it back to none after done debugging. With the option pool-selection, the log will contain pool load balancing algorithm details. This includes common actions taken to a set of pools (for example, whether all pools reset the ratio counter during the algorithm) and the result of the load balancing algorithm (for example, whether a pool is finally selected and the reason if applicable). With the option pool-traversal, the log will contain details of all pools traversed during load balancing. With the option pool-member-selection, the log will contain pool member load balancing algorithm details. This includes common actions taken to a set of pool members and the result of the load balancing algorithm. With the option pool-member-traversal, the log will contain details of all pool members traversed during load balancing. The default value is none.

minimal-response

Specifies GTM will form the smallest allowable DNS response to a query. Typically, this will be a single resource record in the answer section. When set to disabled, GTM will attempt to chase CNAME chains, if required, to obtain the ultimate answer, and it will attempt to add address resource records to the additional section of the response for each answer when needed. The default value is enabled.

metadata

Specifies user-defined data to associate with a server. By default the persist attribute is set to true. This means the data is saved into the configuration file.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

persistence

When enabled, specifies that when a local DNS server makes repetitive requests on behalf of a client, the system reconnects the client to the same resource as previous requests. The default value is disabled.

persist-cidr-ipv4

Specifies a mask used to group IPv4 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

persist-cidr-ipv6

Specifies a mask used to group IPv6 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

pools

Configures the pools the system uses when load balancing requests for this wide IP. The default value is none.

pools-cname

Configures the CNAME pools the system uses when load balancing requests for this wide IP. The default value is none.

pool-lb-mode

Specifies the load balancing method used to select a pool in this wide IP. This option is relevant only when multiple pools are configured for this wide IP. The default value is round-robin.

The available load balancing methods are:

global-availability

Specifies that the system selects a pool by following the order of the Pool list. The system repeatedly selects the first pool in the list for as long as its status is available. If the pool becomes unavailable for any reason, the system then repeatedly selects the next pool in the list, and so on.

ratio

Specifies that the system selects a pool based on the ratio that you assign to the pool.

round-robin

Specifies that the system selects pools sequentially.

topology

Specifies that the system selects a pool based on topology information in the incoming LDNS request. Note that this load balancing method works only if you have configured a topology statement.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rules

Specifies the iRules(r) that this wide IP uses for load balancing decisions. 'when' clauses for each event are grouped across all iRules(r) on this wide IP. For each event, clauses are evaluated in the listed rules order. The default value is none.

ttd-persistence

Specifies, in seconds, the length of time for which a persistence entry is valid. This value can range from 0 through 4294967295 seconds. The default value is 3600.

topology-prefer-edns0-client-subnet

Specifies, when set to enabled, that this wide IP should use the edns0 client subnet option (if one exists) instead of the source address when using topology load balancing. When this option is set to disabled or if the query did not contain a client subnet option, the system will fall back to the source address.

This setting has no effect when the global setting, configured under gtm global-settings load-balancing topology-prefer-edns0-client-subnet [disabled | enabled], is set to enabled. When either setting is enabled then this feature will be enabled.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, reset-stats, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-09-28 gtm wideip mx(1)

gtm wideip naptr

NAME

naptr - Configures a wide IP that accepts NAPTR queries.

MODULE

gtm wideip

SYNTAX

Configure the wideip naptr component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

create wideip naptr [name]

modify wideip naptr [name]

options:

aliases [name...name]

app-service [[string] | none]

description [string]

[disabled | enabled]

failure-rcode [formerr | noerror | notimpl | nxdomain | refused | servfail]

failure-rcode-ttl [integer]

```

failure-rcode-response [disabled | enabled]
last-resort-pool [type name]
load-balancing-decision-log-verbosity [[pool-member-selection | pool-member-traversal | pool-selection | pool-traversal] | none]
minimal-response [disabled | enabled]
metadata none
metadata
  [add | delete | modify | replace-all-with] {
[metadata_name ... ] {
  persist [ true | false ]
  value [ "value content" ]
}
}
persistence [disabled | enabled]
  persist-cidr-ipv4 [integer]
  persist-cidr-ipv6 [integer]
  pool-lb-mode [global-availability | ratio | round-robin | topology]
  pools none
  pools
    [add | delete | modify | replace-all-with] {
[pool name]...
    }
  pools-cname none
  pools-cname
    [add | delete | modify | replace-all-with] {
[pool name]...
    }
  rules none
  rules {
[rule name]
  ...
}
topology-prefer-edns0-client-subnet [disabled | enabled]
ttl-persistence [integer]

```

```
edit wideip naptr [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
  all-properties
  non-default-properties
```

```
reset-stats wideip naptr
```

```
reset-stats wideip naptr [ [name] | [glob] | [regex] ] ... ]
```

```
DISPLAY
```

```
list wideip naptr
```

```
list wideip naptr [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config wideip naptr
```

```
show running-config wideip naptr [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
  all-properties
  non-default-properties
  one-line
  partition
```

```
show wideip naptr
```

```
show wideip naptr [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  (detail | global)
  field-fmt
```

```
DELETE
```

```
delete wideip naptr [all | [name] ]
```

```
DESCRIPTION
```

You can use the wideip component to create, modify, display, or delete a wide IP that responds to NAPTR queries. A NAPTR wide IP is a mapping of a fully-qualified domain name (FQDN) to a set of either services or hosts defined under that name.

```
EXAMPLES
```

```
create wideip naptr example.com
```

Creates a NAPTR wide IP named example.com.

```
delete wideip naptr example.com
```

Deletes the NAPTR wide IP named example.com.

```
OPTIONS
```

```
aliases
```

Specifies alternate domain names for the web site content you are load balancing. You can use two different wildcard characters, asterisk (*) and question mark (?), to represent one or more characters. The default value is none.

```
app-service
```

Specifies the name of the application service to which this wide ip belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this wide ip. Only the application service can modify or delete this wide ip.

description

User defined description. Note: Advanced search on the GUI's Wideip list page can be turned on/off by modifying the DB variable `ui.advancedsearch` via the `tmsh` command "`modify sys db ui.advancedsearch value true/false`". This will result in a new description column and the inclusion of that field in the search.

[disabled | enabled]

Specifies whether the wide IP and its resources are available for load balancing.

`glob` Displays the items that match the `glob` expression. See help `glob` for a description of `glob` expression syntax.

failure-rcode

Specifies the DNS RCODE used when `failure-rcode-response` is enabled. Default is `noerror`. Options include `noerror` (no type exists at this name), `formerr` (format error in query), `servfail` (unable to process query), `nxdomain` (name does not exist), `notimpl` (no support for this kind of query), and `refused` (refuse to process based on policy). If `failure-rcode-ttl` is non-zero, only the Authority section of the `noerror` or `nxdomain` response will include a SOA record.

failure-rcode-response

When enabled, specifies that the system returns a RCODE response to Wide IP requests after exhausting all load-balancing methods. This response is an authoritative empty answer from the system to NAPTR record requests. With this option enabled, the system responds faster to NAPTR requests for which it does not have NAPTR records configured. The default value is `disabled`.

failure-rcode-ttl

Specifies the negative caching TTL of the SOA for the RCODE response. The default is 0, meaning no SOA is included (i.e. no caching).

last-resort-pool

Specifies a last resort pool to use when load balancing requests for this wide IP. Any NAPTR or CNAME pool type is allowed. The default value is `none`.

load-balancing-decision-log-verbosity

Specifies the amount of detail logged when making load balancing decisions. This is used for debugging purpose only. Performance will be affected if the value is not `none`. Please reset it back to `none` after done debugging. With the option `pool-selection`, the log will contain pool load balancing algorithm details. This includes common actions taken to a set of pools (for example, whether all pools reset the ratio counter during the algorithm) and the result of the load balancing algorithm (for example, whether a pool is finally selected and the reason if applicable). With the option `pool-traversal`, the log will contain details of all pools traversed during load balancing. With the option `pool-member-selection`, the log will contain pool member load balancing algorithm details. This includes common actions taken to a set of pool members and the result of the load balancing algorithm. With the option `pool-member-traversal`, the log will contain details of all pool members traversed during load balancing. The default value is `none`.

minimal-response

Specifies GTM will form the smallest allowable DNS response to a query. Typically, this will be a single resource record in the answer section. When set to `disabled`, GTM will attempt to chase CNAME chains, if required, to obtain the ultimate answer, and it will attempt to add address resource records to the additional section of the response for each answer when needed. The default value is `enabled`.

metadata

Specifies user-defined data to associate with a server. By default the `persist` attribute is set to `true`. This means the data is saved into the configuration file.

`name` Specifies a unique name for the component. This option is required for the commands `create` and `modify`.

partition

Displays the administrative partition within which the component resides.

persistence

When enabled, specifies that when a local DNS server makes repetitive requests on behalf of a client, the system reconnects the client to the same resource as previous requests. The default value is `disabled`.

persist-cidr-ipv4

Specifies a mask used to group IPv4 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

persist-cidr-ipv6

Specifies a mask used to group IPv6 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

pools

Configures the pools the system uses when load balancing requests for this wide IP. The default value is `none`.

pools-cname

Configures the CNAME pools the system uses when load balancing requests for this wide IP. The default value is `none`.

pool-lb-mode

Specifies the load balancing method used to select a pool in this wide IP. This option is relevant only when multiple pools are configured for this wide IP. The default value is `round-robin`.

The available load balancing methods are:

`global-availability`

Specifies that the system selects a pool by following the order of the Pool list. The system repeatedly selects the first pool in the list for as long as its status is available. If the pool becomes unavailable for any reason, the system then repeatedly selects the next pool in the list, and so on.

ratio

Specifies that the system selects a pool based on the ratio that you assign to the pool.

round-robin

Specifies that the system selects pools sequentially.

topology

Specifies that the system selects a pool based on topology information in the incoming LDNS request. Note that this load balancing method works only if you have configured a topology statement.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rules

Specifies the iRules(r) that this wide IP uses for load balancing decisions. 'when' clauses for each event are grouped across all iRules(r) on this wide IP. For each event, clauses are evaluated in the listed rules order. The default value is none.

ttl-persistence

Specifies, in seconds, the length of time for which a persistence entry is valid. This value can range from 0 through 4294967295 seconds. The default value is 3600.

topology-prefer-edns0-client-subnet

Specifies, when set to enabled, that this wide IP should use the edns0 client subnet option (if one exists) instead of the source address when using topology load balancing. When this option is set to disabled or if the query did not contain a client subnet option, the system will fall back to the source address.

This setting has no effect when the global setting, configured under gtm global-settings load-balancing topology-prefer-edns0-client-subnet [disabled | enabled], is set to enabled. When either setting is enabled then this feature will be enabled.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-09-28 gtm wideip naptr(1)

gtm wideip srv

NAME

srv - Configures a wide IP that accepts SRV queries.

MODULE

gtm wideip

SYNTAX

Configure the wideip srv component within the gtm module using the syntax shown in the following sections.

CREATE/MODIFY

create wideip srv [name]

modify wideip srv [name]

options:

aliases [name...name]

app-service [[string] | none]

description [string]

[disabled | enabled]

failure-rcode [formerr | noerror | notimpl | nxdomain | refused | servfail]

failure-rcode-ttl [integer]

failure-rcode-response [disabled | enabled]

last-resort-pool [type name]

load-balancing-decision-log-verbosity [[pool-member-selection | pool-member-traversal | pool-selection | pool-traversal] | none]

minimal-response [disabled | enabled]

metadata none

metadata

[add | delete | modify | replace-all-with] {

```

[metadata_name ... ] {
  persist [ true | false ]
  value [ "value content" ]
}
}
persistence [disabled | enabled]
persist-cidr-ipv4 [integer]
persist-cidr-ipv6 [integer]
pool-lb-mode [global-availability | ratio | round-robin | topology]
pools none
pools
  [add | delete | modify | replace-all-with] {
[pool name]...
  }
  pools-cname none
  pools-cname
  [add | delete | modify | replace-all-with] {
[pool name]...
  }
  rules none
  rules {
[rule name]
  ...
  }
  topology-prefer-edns0-client-subnet [disabled | enabled]
  ttl-persistence [integer]

```

```

edit wideip srv [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

```

reset-stats wideip srv
reset-stats wideip srv [ [name] | [glob] | [regex] ] ... ]

```

```

DISPLAY
list wideip srv
list wideip srv [ [name] | [glob] | [regex] ] ... ]
show running-config wideip srv
show running-config wideip srv [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
show wideip srv
show wideip srv [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  (detail | global)
  field-fmt

```

```

DELETE
delete wideip srv [all | [name] ]

```

DESCRIPTION

You can use the wideip component to create, modify, display, or delete a wide IP that responds to SRV queries. An SRV wide IP is a mapping of a fully-qualified domain name (FQDN) specifying a particular service of a domain to a host and port that provides the service.

EXAMPLES

```

create wideip srv _imap._tcp.mydomain.com

Creates an SRV wide IP named _imap._tcp.mydomain.com.

delete wideip srv _imap._tcp.mydomain.com

Deletes the SRV wide IP named _imap._tcp.mydomain.com.

```

OPTIONS

aliases

Specifies alternate domain names for the web site content you are load balancing. You can use two different wildcard characters, asterisk (*) and question mark (?), to represent one or more characters. The default value is none.

app-service

Specifies the name of the application service to which this wide ip belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this wide ip. Only the application service can modify or delete this wide ip.

description

User defined description. Note: Advanced search on the GUI's Wideip list page can be turned on/off by modifying the DB variable ui.advancedsearch via the tmsh command "modify sys db ui.advancedsearch value true/false". This will result in a new description column and the inclusion of that field in the search.

[disabled | enabled]

Specifies whether the wide IP and its resources are available for load balancing.

failure-rcode

Specifies the DNS RCODE used when failure-rcode-response is enabled. Default is noerror. Options include noerror (no type exists at this name), formerr (format error in query), servfail (unable to process query), nxdomain (name does not exist), notimpl (no support for this kind of query), and refused (refuse to process based on policy). If failure-rcode-ttl is non-zero, only the Authority section of the noerror or nxdomain response will include a SOA record.

failure-rcode-response

When enabled, specifies that the system returns a RCODE response to Wide IP requests after exhausting all load-balancing methods. This response is an authoritative empty answer from the system to SRV record requests. With this option enabled, the system responds faster to SRV requests for which it does not have SRV records configured. The default value is disabled.

failure-rcode-ttl

Specifies the negative caching TTL of the SOA for the RCODE response. The default is 0, meaning no SOA is included (i.e. no caching).

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

last-resort-pool

Specifies a last resort pool to use when load balancing requests for this wide IP. Any SRV or CNAME pool type is allowed. The default value is none.

load-balancing-decision-log-verbosity

Specifies the amount of detail logged when making load balancing decisions. This is used for debugging purpose only. Performance will be affected if the value is not none. Please reset it back to none after done debugging. With the option pool-selection, the log will contain pool load balancing algorithm details. This includes common actions taken to a set of pools (for example, whether all pools reset the ratio counter during the algorithm) and the result of the load balancing algorithm (for example, whether a pool is finally selected and the reason if applicable). With the option pool-traversal, the log will contain details of all pools traversed during load balancing. With the option pool-member-selection, the log will contain pool member load balancing algorithm details. This includes common actions taken to a set of pool members and the result of the load balancing algorithm. With the option pool-member-traversal, the log will contain details of all pool members traversed during load balancing. The default value is none.

minimal-response

Specifies GTM will form the smallest allowable DNS response to a query. Typically, this will be a single resource record in the answer section. When set to disabled, GTM will attempt to chase CNAME chains, if required, to obtain the ultimate answer, and it will attempt to add address resource records to the additional section of the response for each answer when needed. The default value is enabled.

metadata

Specifies user-defined data to associate with a server. By default the persist attribute is set to true. This means the data is saved into the configuration file.

`name` Specifies a unique name for the component. The format of an SRV wide IP name is governed by RFC. The general form is '_service._protocol.domain_labels'. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

persistence

When enabled, specifies that when a local DNS server makes repetitive requests on behalf of a client, the system reconnects the client to the same resource as previous requests. The default value is disabled.

persist-cidr-ipv4

Specifies a mask used to group IPv4 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

persist-cidr-ipv6

Specifies a mask used to group IPv6 LDNS addresses. This feature allows one persistence record to be shared by LDNS addresses that match within this mask.

pools

Configures the pools the system uses when load balancing requests for this wide IP. The default value is none.

pools-cname

Configures the CNAME pools the system uses when load balancing requests for this wide IP. The default value is none.

pool-lb-mode

Specifies the load balancing method used to select a pool in this wide IP. This option is relevant only when multiple pools are configured for this wide IP. The default value is round-robin.

The available load balancing methods are:

global-availability

Specifies that the system selects a pool by following the order of the Pool list. The system repeatedly selects the first pool in the list for as long as its status is available. If the pool becomes unavailable for any reason, the system then repeatedly selects the next pool in the list, and so on.

ratio

Specifies that the system selects a pool based on the ratio that you assign to the pool.

round-robin

Specifies that the system selects pools sequentially.

topology

Specifies that the system selects a pool based on topology information in the incoming LDNS request. Note that this load balancing method works only if you have configured a topology statement.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rules

Specifies the iRules(r) that this wide IP uses for load balancing decisions. 'when' clauses for each event are grouped across all iRules(r) on this wide IP. For each event, clauses are evaluated in the listed rules order. The default value is none.

ttd-persistence

Specifies, in seconds, the length of time for which a persistence entry is valid. This value can range from 0 through 4294967295 seconds. The default value is 3600.

topology-prefer-edns0-client-subnet

Specifies, when set to enabled, that this wide IP should use the edns0 client subnet option (if one exists) instead of the source address when using topology load balancing. When this option is set to disabled or if the query did not contain a client subnet option, the system will fall back to the source address.

This setting has no effect when the global setting, configured under gtm global-settings load-balancing topology-prefer-edns0-client-subnet [disabled | enabled], is set to enabled. When either setting is enabled then this feature will be enabled.

SEE ALSO

create, delete, edit, glob, gtm pool, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-09-28 gtm wideip srv(1)

ilx

ilx global-settings

NAME

global-settings - system-wide settings for iRules Language Extension (ILX) Plugin.

MODULE

ilx

SYNTAX

The global-settings can be displayed/modified with the following syntax:

MODIFY

modify global-settings [name]

options:

debug-port-blacklist add/delete/replace-all-with { [port] }

log-publisher [name]

DISPLAY

list global-settings [name]

options:

debug-port-blacklist

DESCRIPTION

An iRules Language Extension (ILX) plugin global-settings profile a way to store system-wide settings.

EXAMPLES

list ilx global-settings

See the current settings

modify global-settings debug-port-blacklist add { 20001 }

Add port 20001 to the list of reserved ports that cannot be used to attach node inspector.

OPTIONS

debug-port-blacklist

Specifies the list of reserved ports that cannot be used to attach node inspector in case debugging is enabled.

log-publisher Specifies the log destination for messages generated by the ILXLogger nodejs object in an extension without a plugin or extension level log publisher.

SEE ALSO

ilx plugin, modify, list

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-03-10 ilx global-settings(1)

ilx node-version

NAME

node-version - Available Node.js versions for iRules Language Extension (ILX) Plugin and Workspace.

MODULE

ilx

SYNTAX

The node-version can be displayed with the following syntax:

DISPLAY

list node-version [name]

DESCRIPTION

Provides major and full version of available node.js versions available to an iRules Language Extension (ILX) Plugin and Workspace.

EXAMPLES

list ilx node-version

See the current available versions

SEE ALSO

ilx plugin, ilx workspace, list

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-09-20 ilx node-version(1)

ilx plugin

NAME

plugin - Configures an iRules Language Extension (ILX) Plugin.

MODULE

ilx

SYNTAX

Configure the plugin component within the ilx module using the syntax shown in the following sections.

CREATE/MODIFY

create plugin [name]

options:

app-service [name]

description [string]

disabled

enabled
from-workspace [name]
log-publisher [name]

modify plugin [name]
options:
app-service [name]
description [string]
disabled
enabled
extensions {
 [name] {
 command-arguments [list of strings]
 command-options [list of strings]
 concurrency-mode [dedicated | single]
 data-groups [[[name] ...] | none]
 debug-port-range-high [port]
 debug-port-range-low [port]
 description [string]
 heartbeat-interval [seconds]
 ilx-logging [disabled | enabled]
 log-publisher [name]
 max-restarts [integer]
 restart-interval [integer]
 trace-level [integer]
 }
}

from-workspace [name]
log-publisher [name]

reset-stats plugin [[[name] | [glob] | [regex]] ...]
options:
extensions { [name] ... }

DISPLAY
list plugin
list plugin [[[name] | [glob] | [regex]] ...]
show running-config plugin
show running-config plugin [[[name] | [glob] | [regex]] ...]
options:
all
all-properties
disk-space
node-version
non-default-properties
one-line
partition
recursive
staged-directory

show plugin
show plugin [[name] ...]
options:
all
all-properties (default | exa | gig | kil | meg | peta | raw | tera |
 yotta | zetta)
extensions { [name] ... }
field-fmt
recursive

DELETE
delete plugin [name]
options:
all
recursive

START
start plugin [name]

STOP/RESTART
restart plugin [name]
stop plugin [name]
options:
drain-existing-connections
immediate
stop-draining

DESCRIPTION

An iRules Language Extension (ILX) plugin provides the following functionality:
An RPC interface that allows an iRule to communicate with an ILX Node.js(tm) process.
A Streaming interface that allows an ILX Node.js(tm) process to manage Virtual Server connections and payload data.

A plugin is created from an ilx workspace. Files from the workspace are copied into a system area and run from the system area. iRule files contained in the workspace are read into the system as ltm rule configuration. The ltm rule configuration objects are created in a sys folder that has the same name as the

plugin.

The system will start one or more Node.js(tm) process for each extension (see also concurrency-mode below). Node.js(tm) will look in the extension directory for the file package.json. Node.js(tm) will look in package.json for a main field that identifies the main entry point of the plugin. If the main field is not present node will look for the file index.js.

The status of plugin processes can be viewed with the show command.

```
show plugin my_plugin
```

A plugin can be started, stopped and restarted. When a plugin is initially created the system will start processes for each extension. When a plugin is stopped existing plugin processes associated with active tmm connections are allowed to drain. Draining processes continue to run until the connections are closed. There are several options that control how processes are stopped, drain-existing-connections, immediate and stop-draining, and are describe below.

There are several ILX commands that an iRule may use to communicate with a Node.js(tm) extension process: ILX::init, ILX::call and ILX::notify.

ILX::init [plugin name] [extension name] establishes a communication path from an iRule to the node process. It returns a handle that must be passed to ILX::call and ILX::notify.

ILX::call [ILX handle] [method] [optional arg 1] [optional arg 2] invokes the specified node method in the extension associated with the ILX handle. If the node method returns a javascript array or dictionary then a Tcl list object is returned from ILX::call, otherwise a Tcl string object is returned.

ILX::notify [ILX handle] [method] [optional arg 1] [optional arg 2] sends a message to the node method but does not wait for the method to complete before returning. TCL_OK is returned if the message was successfully queued. Any other return value indicates the message could not be queued.

RPC

The RPC API consists of two primary components, Node.js and Tcl.

The Node.js API contains an ILXServer class. Methods are added to the server using ILXServer.addMethod("method name", function(req, res) { ... }). "req" is an ILXRequest object and "res" is an ILXResponse object. ILXServer.listen will start the server.

A Tcl iRule connects to the Node.js process using the Tcl command
set plugin [ILX::init "plugin name" "extension name"]

The Javascript methods are invoked from a Tcl iRule using either ILX::call or ILX::notify. The args are user defined and passed to the Node.js method via the ILXRequest object.
set response [ILX::call \$plugin "method name" args...]

ILX::notify is similar to ILX::call except that the iRule will not wait for a response from the Node.js process.

```
ILX::call $plugin "method name" args...]
```

Here is a Hello World Node.js script

```
var f5 = require('f5-nodejs');
var server = new f5.ILXServer();
server.addMethod("hello_world", function(req, res) {
  console.log("request args: " + JSON.stringify(req));
  res.reply('hello_world');
});
server.listen();
```

Here is the Hello World Tcl iRule that passes some data to the Node.js process, waits for a response and logs the response to /var/log/ltm. The iRule must be associated with an ltm virtual-server.

```
when CLIENT_ACCEPTED {
  set plugin [ILX::init rpc example]
  set rv [ILX::call $plugin hello_world arg1 val1 arg2 val2]
  log local0.debug $rv
}
```

Run the following tmsh commands to create the ilx workspace and ilx plugin configuration for the RPC example above. See also help ilx workspace for more examples describing how to manage a workspace.

```
create ilx workspace rpc
```

```
create ilx workspace rpc extension example
```

```
edit ilx workspace rpc extension example
```

A file named index.js will be created. Replace the content of index.js with the Node.js Hello World script above. If you have bash shell access you can directly edit the file /var/ilx/workspaces/Common/rpc/extensions/example/index.js.

```
edit ilx workspace rpc rule example
```

A file named rpc_example.tcl will be created. Replace the content of rpc_example.tcl with the Tcl iRule Hello World example above. If you have bash shell access you can directly edit the file

```
/var/ilx/workspaces/Common/rpc/rules/hello_world.tcl.
```

```
create ilx plugin rpc from-workspace rpc  
modify ltm virtual [name] rules { rpc/example }
```

You must create a Virtual Server and associate the `rpc_example` rule with the Virtual Server. When a client connects to the Virtual Server you should see the output from both the Node.js process and Tcl iRule in `/var/log/ltm`.

See also the BIG-IP iRulesLX User Guide located in the Programmability Wiki at devcentral.f5.com.

Streaming

The Streaming API allows a Node.js Plugin process to act as a proxy for connections through a Virtual Server. When a client connects to a Virtual Server an ILXFlow object is created and passed to the plugin. The ILXFlow object contains two Node.js sockets. `ILXFlow.client` represents the connection from the client to the plugin. `ILXFlow.server` represents the connection from the plugin to a server. The plugin will use `ILXFlow.client` to receive or send data to the client. The plugin will use `ILXFlow.server` to send or receive data from the server. The plugin may inspect or modify any/all client or server data. The `ILXFlow.client` and `ILXFlow.server` objects behave as native Node.js sockets. The plugin can also be configured to act as a server and respond directly to the client without completing a connection to a server. The plugin can also be configured as a native Node.js HTTP server.

```
[Client] <-> [Virtual Server] <-> [TMM]  
^  
|  
v  
[ILXFlow.client]  
[Plugin]  
[ILXFlow.server]  
^  
|  
v  
[TMM] <-> [Server]
```

Note that the plugin sees TCP payload. TCP headers are stripped by the TMM and are not visible to the plugin.

The following is a brief overview of the Javascript classes that make up the Streaming API. These classes contain many more features and options as described in the BIG-IP iRulesLX User Guide located in the Programmability Wiki at devcentral.f5.com.

ILXPlugin - Connects to the TMM to accept Virtual Server client connections. When a connection is made an ILXFlow object is passed to the plugin via the "connect" event: `plugin.on("connect", function(flow) { ... })`.

ILXPluginOptions - Options that configure the events and data that a plugin will received.

ILXFlow - An object that represents a client socket to Virtual Server and server socket to a server.

ILXStream - An object that represents the client side connection from a client to a Virtual Server (`ILXFlow.client`), the server side connection from a Virtual Server to a server (`ILXFlow.server`), or a plugin initiated connection to a server (created using `ILXStream.connect`).

ILXLbOptions - Options that allow the plugin to specify which server the client will connect to by calling `ILXFlow.lbSelect(ILXLbOptions)`.

ILXTable - Functions for reading and writing the TMM Session DB. Very similar interface to the Tcl iRule table command.

ILXBufferUtil - Functions for searching and modifying a Node.js Buffer object.

ILXDatagroup - Functions for accessing ltm data-group configuration. The set of data-groups available to a plugin must be specified using the `ILXPluginOptions` when the plugin is started.

ILXTransaction - If HTTP and ILX profiles are attached to a virtual server then all header and data inspection and modification are done in the context of request and response transaction. An `ILXTransaction` object is passed to the following events: `requestStart`, `requestComplete`, `responseStart` and `responseComplete`. The `ILXTransaction` object contains things like request headers, URI and method, and response headers and status.

The following is a sample streaming plugin that acts as a pass through. Data is read from the client and sent to the server. Data is read from the server and sent to the client. Note that it is important to handle "error" events as the Node.js plugin process will exit if they are not handled.

```
var f5 = require("f5-nodejs");  
var plugin = new f5.ILXPlugin();  
plugin.on("connect", function(flow) {  
  flow.client.on("readable", function() {  
    var buffer;  
    while (true) {  
      buffer = flow.client.read();  
      if (buffer === null) {  
        break;  
      }  
      console.log("Client Data");  
      console.log(buffer.toString());  
      flow.server.write(buffer);  
    }  
  }  
})
```

```

});
flow.client.on("error", function(err) {
console.log("client socket error: " + err);
});
flow.server.on("readable", function() {
var buffer;
while (true) {
buffer = flow.server.read();
if (buffer === null) {
break;
}
console.log("Server Data");
console.log(buffer.toString());
flow.client.write(buffer);
}
});
flow.server.on("error", function(err) {
console.log("server socket error: " + err);
});
flow.on("error", function(err) {
console.log("flow error: " + err);
});
});
var options = new f5.ILXPluginOptions(); // use defaults
plugin.start(options);

```

The following is a sample streaming HTTP plugin that acts as a pass through. Note that the example above can also be used to manage HTTP traffic. This example requires a Virtual Server with both HTTP and ILX profiles.

```

var f5 = require('f5-nodejs');
var plugin = new f5.ILXPlugin();
plugin.on("connect", function(flow)
{
flow.client.on("requestStart", function(request) {
// HTTP request URI, headers and more
console.log(
"requestStart: " + JSON.stringify(request.params));
});
flow.client.on("readable", function() {
// HTTP request body, if present
var buf;
while (true) {
buf = flow.client.read();
if (buf !== null) {
console.log("request body:" + buf.length + " bytes");
console.log(buf.toString());
flow.server.write(buf);
}
else {
break;
}
}
});
flow.client.on("requestComplete", function(request) {
// Entire HTTP request received
console.log(
"requestComplete: " + JSON.stringify(request.params));
request.complete();
});
flow.client.on("error", function(err) {
console.error("flow.client error:" + err);
});

flow.server.on("responseStart", function(response) {
// HTTP response headers, status and more
console.log(
"responseStart: " + JSON.stringify(response.params));
});
flow.server.on("readable", function() {
// HTTP response body, if present
var buf;
while (true) {
buf = flow.server.read();
if (buf !== null) {
console.log(
"response body:" + buf.length + " bytes");
console.log(buf.toString());
flow.client.write(buf);
}
else {
break;
}
}
});
flow.server.on("responseComplete", function(response) {
// Entire HTTP response received
console.log(

```

```

    "responseComplete: " + JSON.stringify(response.params));
response.complete();
});
    flow.server.on("error", function(err) {
console.error("flow.server error:" + err);
    });
    flow.on("error", function(err) {
console.error("flow.error:" + err);
    });
});
var options = new f5.ILXPluginOptions(); // use defaults
plugin.start(options);

```

Run the following tmsh commands to create ilx workspace and ilx plugin configuration for a streaming plugin using either of the sample plugins above.

```
create ilx workspace streaming
```

```
create ilx workspace streaming extension example
```

```
edit ilx workspace streaming extension example
```

A file named index.js will be created. Replace the content of index.js with the Node.js streaming example from above. If you have bash shell access you can directly edit the file `/var/ilx/workspaces/Common/streaming/extensions/example/index.js`.

```
create ilx plugin streaming from-workspace streaming
```

```
modify ilx plugin streaming extensions { example { ilx-logging
enabled concurrency-mode single } }
```

When ilx-logging is enabled all stdout and stderr from the plugin is sent to `/var/log/ilx/...`, in this case `/var/log/ilx/Common.streaming.example`.

```
create ltm profile ilx streaming-profile plugin streaming
```

```
modify ltm virtual [name] profiles replace-all-with {
tcp streaming-profile }
```

You must create a Virtual Server and associate the streaming-profile with the Virtual Server. When a client connects to the Virtual Server the client should receive the server data. The plugin also logs the client and server data that is received to `/var/log/ilx/Common.streaming.example`. If you are using the HTTP sample you will need to assign an HTTP profile to the virtual server.

EXAMPLES

```
create plugin my_plugin from-workspace my_workspace
```

Creates a plugin named my_plugin. Rule files from the workspace are read into the system configuration and placed in the sys folder my_plugin.

```
modify plugin my_plugin from-workspace my_workspace
```

Restart the plugin using the current files in my_workspace.

```
modify plugin my_plugin extensions { my_extension { concurrency-mode dedicated } }
```

Modify the concurrency-mode for a specific extension.

```
delete plugin my_plugin
```

Deletes the plugin named my_plugin. The iRule configuration in the sys folder named my_plugin will also be deleted. The sys folder my_plugin will be deleted if there are no remaining (non-plugin related) configuration items in the folder.

```
list plugin
```

List the configuration for all plugins in the current sys folder.

```
list plugin my_plugin
```

List the configuration for my_plugin.

```
show plugin my_plugin
```

Displays the status of all extension processes running as part of my_plugin.

```
show plugin my_plugin extensions { my_extension }
```

Displays the status of processes in extension my_extension running as part of my_plugin.

```
reset-stats plugin my_plugin
```

Reset all statistics for my_plugin.

```
reset-stats plugin my_plugin extensions { my_extension }
```

Reset statistics for the extension my_extension in my_plugin.

```
start plugin my_plugin
```

Starts all extensions in my_plugin.

```
stop plugin my_plugin
```

Stops all extensions in my_plugin. Active processes and connections are allowed to drain.

```
stop plugin my_plugin drain-existing-connections
```

Stops all extension processes in my_plugin. Active connections that are currently using the plugin are allowed to drain. The plugin processes will be stopped when all associated active connections are closed.

```
stop plugin my_plugin immediate
```

Immediately stops all extension processes associated with my_plugin. New and existing connections to Virtual Servers that have iRules that interact with the plugin will receive a Tcl error when using the Tcl functions ILX::init and ILX::call.

```
stop plugin my_plugin stop-draining
```

Immediately stops all extension processes in my_plugin that are draining. Existing connections that have iRules that interact with the plugin and were attached to the draining processes will receive a Tcl error when using the Tcl functions ILX::init and ILX::call.

OPTIONS

app-service

Specifies the name of the application service to which the pool belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the plugin. Only the application service can modify or delete the plugin.

description

User defined description.

disabled

Disable the plugin. All node processes associated with the plugin that are associated with active tmm connections are allowed to continue until the connections close. To immediately stop all processes use the stop command.

enabled

Enable the plugin. Node processes will be started for each plugin extension.

extensions

Modify attributes of plugin extensions.

command-arguments

The optional list of arguments to be passed to the nodejs (after the extension).

Eg.:

```
command-arguments { arg1 arg2 }
```

will result in nodejs process started like this:

```
/usr/bin/nodejs /extensions/ arg1 arg2
```

The plugin must be restarted for this setting to take affect.

command-options

The list of options passed to Node.js before the name of the extension script. These must be valid Node.js options. The Node.js process will fail to start if it is provided with invalid options.

Example:

```
command-options { --debug-brk=12345 }
```

will result in nodejs process started like this:

```
/usr/bin/nodejs --debug-brk=12345 /extensions/
```

Please note the choice of the specified port is not guaranteed. If the specified port is not available, the system will choose another consecutive port (see Debugging section for more details).

The plugin must be restarted for this setting to take affect.

concurrency-mode

Specifies a concurrency mode for the extension.

The value dedicated specifies that a node process should be started for each tmm that has been provisioned (see sys db provision.tmmcountactual). Each node process will be dedicated to servicing requests from a single tmm.

The value single specifies that a single node process should be started. This node process will service requests from all tmm processes in the system.

data-groups

The set of data-groups that will be made available to the plugin through the ILXDatagroup JavaScript class. The default value is none.

debug-port-range-high

When debugging is enabled and the debug port is not specified in command-options, ILX will assign the port, searching for an available port to attach the node inspector, see the section on Debugging. It will start searching at the debug-port-range-low value and try each port until it reaches the debug-port-range-high value.

debug-port-range-low

When debugging is enabled and the debug port is not specified in command-options, ILX will assign the port, searching for an available port to attach the node inspector, see the section on Debugging. It will start searching at the debug-port-range-low value and try each port until it reaches the debug-port-range-high value.

description

User defined description.

heartbeat-interval

Specifies the maximum number of seconds between extension process heartbeats. The extension process heartbeat is updated by the BIG-IP version of libuv. The BIG-IP monitors the heartbeat of each extension process. When enabled the system will restart an extension process if the extension does not update its heartbeat over the configured interval. A value of 0 (zero) disables heartbeat checking. The heartbeat should be disabled when the extension is being debugged.

ilx-logging

Specifies where standard out and standard error generated by the plugin will be logged. A value of disabled will cause output to be logged to /var/log/ltm. A value of enabled will cause output to be logged to /var/log/ilx/[plugin].[extension] (enabled). A CPU identifier will be appended to the log file name if the concurrency-mode is dedicated. Setting the value to enabled is intended to aid in debugging a plugin and not overrun /var/log/ltm with debug logs. The default value is disabled.

log-publisher

Specifies the log destination for messages generated by the ILXLogger nodejs object in an extension. See sys log-config publisher.

node-version

Display the Node.js version that is used to run the plugin.

max-restarts

Specifies the maximum number of times an extension process may fail before the system will no longer automatically restart the process. See also restart-interval. The default value is 5.

restart-interval

Specifies the interval, in seconds, over which max-restarts may occur. The default value is 60 seconds.

trace-level

Specifies the global trace level for an extension. This can also be controlled by an extension using the javascript function ILXPlugin.setGlobalTraceLevel (note that this does not update the extension configuration, only the running extension). An extension may use this setting in Javascript by calling ILXPlugin.globalTraceLevel(). F5 extension internals also use this setting to log interaction between the configuration system, the TMM and the plugin extension. A value of zero will stop trace logging. A value of 1 will log configuration messages. A value of 10 will log general function call location information. A value of 20 will include function call information in the packet path. A value of 30 will include packet content that is sent and received to and from the TMM. Note that if the trace level has been set by javascript code in the plugin using ILXPlugin.setGlobalTraceLevel and the extension trace-level setting is zero then to disable tracing in the plugin you must update the plugin code. See also ilx-logging.

from-workspace

Create or modify a plugin using the set of files in the specified workspace. The workspace files are copied into a system are. New node process are started using the latest copy of workspace files. Existing node processes continue to run until all tmm connections associated with those processes have closed.

log-publisher

Specifies the log destination for messages generated by the ILXLogger nodejs object without an extension level log-publisher. See sys log-config publisher.

name Specifies a unique name for the component. This option is required for the commands create and modify.

staged-directory

Display the workspace directory that was used to create the plugin.

Debugging

In order to debug an ILX extension the user must enable debugging via command-options and restart the plugin. Node.js provides three options for enabling remote debugging: --debug[=], --debug-port= and --debug-brk=.

If a debug port is specified the system will search for an available port starting with the specified port. Please note that the debugging can only be enabled in single concurrency mode.

```
modify ilx plugin my_plugin extension { my_extension { command_options add { --debug } } }  
restart ilx plugin my_plugin
```

If a debug port is not specified with either `--debug=`, `--debug-brk=` or `--debug-port=`, the system will search for an available port within the configured `debug-port-range-low` and `debug-port-range-high` range. If a port is specified the system will perform a search for an available port using the specified port instead of the `debug-port-range-low` value. The assignment of the specified port is best-effort and is not guaranteed. In both cases the search for ports is consecutive.

There is a system-wide list of ports that will be skipped during the debug port search:

```
Eg.
# list ilx global-settings
ilx global-settings {
  debug-port-blacklist { 47019 54321 60000 }
}
```

The `node-inspector` process must be started at the bash command line, specifying a valid IP address that is either a self-ip (`net self-ip`) or the management IP address (`sys management-ip`).

```
/usr/local/node/v6/lib/node_modules/.bin/node-inspector --web-host --no-inject
```

It will then be possible to connect to the IP address using the port that the system chose. The debug port that the system chose can be seen by running the command `show plugin my_plugin`.

```
show plugin my_plugin extensions { my_extension }
```

```
...
-----
| Extension Process: my_extension
-----
| Status    running
| PID      1190
| TMM      0
| CPU Utilization (%) 0
| Debug Port 20000
| Memory (bytes)
| Total Virtual Size 0
| Resident Set Size 0
| RPC Info
...

```

It is then possible to connect with a suitable debugging client, such as the Chrome(tm) browser via:
`https://:8080/debug?port=20000`

The externally reachable port opened by `node-inspector` is `":8080"`. Port `"20000"` is the port displayed by the `show ilx plugin` command.

Statistics

Provided the ILX-RPC and ILX Streaming can be used simultaneously there are designated sections for each in the stats.

The plugin statistics are reset every time the plugin is restarted. Additionally the statistics can be reset from the `tms` (see `EXAMPLES` above).

The output looks like the following:

```
-----
ILX::Plugin: my_plugin
-----
State enabled
Log Publisher local-db-publisher

-----
| Extension: my_extension
-----
| Status    running
| CPU Utilization (%) 0
| Memory (bytes)
| Total Virtual Size 0
| Resident Set Size 0
| RPC Info    - ILX-RPC Specific information
| RPC Connections
|   Active    0 - active RPC connections
|   Total    0 - total RPC connections (ILX::init calls)
| Total Calls 0 - total RPC calls (ILX::call calls)
| Notifies    0
| Timeouts   0 - number of expired connections
| Errors     0 - number of connection errors
| Octets In  0 - number of octets received
| Octets Out 0 - number of octets sent
| Average Latency 0 - average time that it takes node to respond to an iRule ILX::call
| Max Latency 0 - max response time it took node to respond to an iRule ILX::call
| Streaming Info - ILX Streaming specific information
| Flow Connections - virtual server initiated connections.
| Clientside
|   Active    0 - active clientside connections
|   Total    0 - total connections
| Serverside
|   Active    0 - active serverside connections
|   Total    0 - total connections

```

```
| Sideband Connections - connections initiated with .connect() call
|   Active    0 - active connections
|   Total    0 - total connections
| Aborts     0 - number of aborted connections
| Timeouts   0 - number of expired connections
| Errors     0 - number of connection errors
| Octets In  0 - number of octets received
| Octets Out  0 - number of octets sent
| Table Calls 0 - number of API calls to SessionDB (ILXTable)
| Datagroup Calls 0 - number of API calls to Datagroups (ILXDatagroup)
| Restarts   0 - number of plugin restarts
| Failures   0 - number of plugin failures
```

```
-----
| Extension Process: my_extension
-----
```

```
| Status running
| PID 20626
| TMM 0
| CPU Utilization (%) 0
| Debug Port 20001
| Memory (bytes)
|   Total Virtual Size 735.6M
|   Resident Set Size 15.2K
| RPC Info
|   MPI channel mem://ilx:1:0 - communication channel between node.js process and TMM
|   RPC Connections
|   ...
|   Process specific fields for ILX-RPC similar to those in the Extension output
|   ...
| Streaming Info
|   MPI channel mem://ilx:2:0 - communication channel between node.js process and TMM
|   Flow Connections
|   ...
|   Process specific fields for ILX-Streaming similar to those in the Extension output
|   ...
```

```
-----
| Extension Process: my_extension
-----
```

```
| Status running
| PID 20627
| TMM 1
| CPU Utilization (%) 0
| ...
```

For ILX-Streaming, more information is available from the ILX profile statistics.

SEE ALSO

create, delete, edit, ilx workspace, list, ltm data-group, ltm virtual-server, modify, sys folder, sys log-config publisher, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-04-06 ilx plugin(1)

ilx workspace

NAME

workspace - Configures a development workspace for Node.js(tm) iRule Language Extensions.

MODULE

ilx

SYNTAX

Configure the workspace component within the ilx module using the syntax shown in the following sections.

CREATE/MODIFY/EDIT

```
create workspace [name]
```

options:

archive [name]

extension [name]

file [name]

from-archive [name]

from-plugin [name]
from-uri [uri]
from-workspace [name]
node-version [version]
rule [name]

modify workspace [name]
options:
 from-archive [name]
 from-plugin [name]
 from-uri [uri]
 from-workspace [name]
 node-version [version]

edit workspace [name]
options:
 extension [name]
 file [name]
 rule [name]

DISPLAY

list workspace
list workspace [[name] ...]
show running-config workspace
show running-config workspace [[name] ...]
options:
 all-files
 all-properties
 archive [name]
 extension [name]
 file [name]
 node-version
 one-line
 partition
 recursive
 rule [name]
 staged-directory
 version

show workspace
show workspace [name]
options:
 all-properties
 detail
 field-fmt

DELETE

delete workspace [name]
options:
 archive [name]
 extension [name]
 file [name]
 rule [name]

DESCRIPTION

An ILX Workspace is an area of the file system used to store Node.js(tm) and supporting files used to construct an ilx plugin. Workspaces are stored in /var/ilx/workspaces/.../.

A workspace contains several components.

/var/ilx/workspaces/Common/extensions/...
/var/ilx/workspaces/Common/rules/...
/var/ilx/workspaces/Common/version

The extensions directory contains sub-directories. Each sub-directory contains files for a specific extension. /var/ilx/workspaces/Common/extensions//... Each extension directory must contain a package.json file. This is a standard npm package definition file. For more information see www.npmjs.org/doc/package.json.html. When an ilx plugin is started node will look in package.json for a main field that identifies the main entry point of the plugin. If the main field is not present node will look for the file index.js.

The rules directory contains zero or more files that contain iRules (see ltm rule). /var/ilx/workspaces/Common/rules/.tcl. ILX iRule commands are used to communicate with a Node.js(tm) process that is running as part of an ILX Plugin. When a new rule is created using either of the following commands a sample iRule file is opened with syntax examples for the ILX iRule commands.

```
create ilx workspace ilx rule test  
edit ilx workspace ilx rule test
```

When a new extension is created using the following commands a sample index.js file will be created that contains Node.js(tm) and ILX iRule examples. The extension directory is also populated with the file package.json and locally installed Node.js(tm) modules.

```
tmsh create ilx workspace my_workspace  
tmsh create ilx workspace my_workspace extension my_extension
```

The version file indicates the BIG-IP version used to create the workspace.

EXAMPLES

create workspace my_workspace

Creates a workspace named my_workspace in /var/ilx/workspaces/Common/my_workspace (assuming your current working folder is the default BIG-IP folder /Common).

create workspace my_workspace rule my_rule

Adds the iRule my_rule as the file my_rule.tcl to the workspace directory named /var/ilx/workspaces/Common/my_workspace/rules/.

edit workspace my_workspace rule my_rule

Adds or modifies the iRule my_rule as the file my_rule.tcl to the workspace directory named /var/ilx/workspaces/Common/my_workspace/rules/.

create workspace my_workspace extension my_extension

Adds a new extension directory in /var/ilx/workspaces/Common/my_workspace/extensions/my_extension and installs node_modules/f5-nodejs within the directory. my_extension must be a valid Linux directory name.

edit workspace my_workspace extension my_extension

If the extension does not exist a new extension directory is created in /var/ilx/workspaces/Common/my_workspace/extensions/my_extension, installs node_modules/f5-nodejs within the directory and opens the editor on the file within the directory named index.js. If the extension already exists the editor will open index.js.

edit workspace my_workspace file extensions/my_extension/my_file.js

If the extension does not exist a new extension directory is created in /var/ilx/workspaces/Common/my_workspace/extensions/my_extension, installs node_modules/f5-nodejs within the directory and opens the editor on my_file.js. If the extension already exists the editor will open my_file.js.

edit workspace my_workspace file extensions/my_extension/my_file.js

The editor is opened on /var/ilx/workspaces/Common/my_workspace/extensions/my_extension/my_file.js. The extension must already exist.

delete workspace my_workspace

Deletes the workspace named my_workspace. This removes all disk files associated with the workspace and cannot be undone. You will be prompted to confirm before the workspace is deleted.

delete workspace my_workspace extension my_extension

Deletes the extension named my_extension in my_workspace.

delete workspace my_workspace rule my_rule

Deletes the iRule file named my_rule in my_extension in my_workspace.

delete workspace my_workspace file extensions/my_extension/my_file

Deletes the file named my_file in my_extension in my_workspace.

list workspace my_workspace

List the content of my_workspace

list workspace recursive my_workspace

List all files and directories in my_workspace.

list workspace my_workspace rule

List the rules associated with my_workspace.

list workspace my_workspace rule my_rule

List the text content of the rule my_rule associated with my_workspace

list workspace my_workspace file rules/my_rule.tcl

List the text content of the file rules/my_rule.tcl associated with my_workspace

list workspace my_workspace extension

List the extensions associated with my_workspace.

list workspace my_workspace extension my_extension

List the content of the my_extension extension directory associated with my_workspace.

list workspace my_workspace file extensions/my_extension/my_file

List the text content of the file my_file in extension my_extension.

show workspace my_workspace

Show the top level files for the specified workspace. Show the group, last modify time, owner, permissions and size of the workspace.

show workspace

Show all workspaces.

OPTIONS

archive

Specifies the name of a workspace archive file. See the section below on archives.

all-files

Recursively list all directory and file names contained in a workspace.

all-properties

For a show command, show group, last modify time, owner, permissions, type and size of the file/directory.

detail

For a show command, show all the files in the workspace.

extension

Specifies the name of a workspace extension.

file For the edit and list commands the file option specifies the name of a file in a workspace to edit or list. For the delete command the file option specifies the name of a workspace file to delete.

from-archive

Specifies the name of a tar or gzipped tar archive file from which to create or update a workspace. See the section below on archives.

from-plugin

Specifies the name of an ilx plugin from which to create or modify a workspace. The files used to create the plugin are copied into the workspace.

from-uri

Specifies the name of a URI from which to create or update a workspace. The URI must refer to a tar or gzipped tar archive file.

from-workspace

Specifies the name of a workspace from which to create or modify a workspace. The files used to create the workspace specified by from-workspace are copied into the new workspace.

node-version

Specifies the version of Node.js that will be used to run an . When an is created or updated from a workspace the plugin inherits the node-version from the workspace. The node-version may be specified when an ilx workspace is created. The node-version may also be changed after the ilx workspace has been created.

When the node-version of a workspace is changed the workspace is archived. The following files and directories are replaced: The file

[workspace name]/extensions/[extension name]/node_modules/f5-nodejs/package.json

and the content of the directory

[workspace name]/extensions/[extension name]/node_modules/f5-nodejs/lib.

To update an existing ilx plugin to use the new workspace version issue the following command:

modify ilx plugin [plugin name] from-workspace [workspace name]

When a node process is started the BigIP sets the NODE_PATH environment variable. For Node.js version 0.12 plugins the NODE_PATH is set to /usr/local/node/v0.12/lib/node_modules. For Node.js version 6 the NODE_PATH is set to /usr/lib/node_modules:/usr/local/node/v6/lib/node_modules.

one-line

List the content of a workspace on a single line.

recursive

Recursively list all workspaces contained in the current BIG-IP system folder (see sys folder).

rule Specifies the name of a rule file to create, delete, edit or list.

staged-directory

The directory where the workspace and associated files reside. This is a read-only property.

name Specifies the name of a workspace. This option is required for the commands create, delete, edit and modify.

ACCESS CONTROL

Workspaces are stored under file system directory /var/ilx/workspaces.

A user with the role of admin can access all workspaces under /var/ilx/workspaces.

Users with the role of irule-manager, manager and resource-admin can manage workspaces under /var/ilx/workspaces/, where is the name of an existing sys folder that the user has permission to access, see also help sys folder and help pwd.

As a convenience workspaces are automatically created in the users' current folder relative to /var/ilx/workspaces. However, workspaces may be created in file system directories that do not correspond to BIG-IP system configuration folders. In this case the workspace must be referred to by full path.

The following creates two workspaces, the first in the file system location /var/ilx/workspaces/Common/my_workspace and the second in the filesystem location /var/ilx/workspaces/Common/my_folder/my_workspace.

```
(tmos)# pwd
/Common
(tmos)# create ilx workspace my_workspace
(tmos)# cd my_folder
(tmos)# pwd
/Common/my_folder
(tmos)# create ilx workspace my_workspace_2
(tmos)# list ilx workspace my_workspace_2

(tmos)# cd ..
(tmos)# pwd
/Common
(tmos)# list ilx workspace my_workspace
```

Full folder paths may be used to refer to a workspace.

```
(tmos)# create ilx workspace /Common/my_folder/my_workspace_2
(tmos)# list ilx workspace /Common/my_folder/my_workspace_2
```

If an auth partition is removed from the BIG-IP configuration the workspace files remain on disk and can be accessed by a user with the role of admin. If a sys folder is removed under an auth partition, such as removing the folder /Common/my_config, all workspaces under the file system directory /var/ilx/workspaces/Common/my_config remain accessible to all users that have access to the Common partition.

ARCHIVES

A workspace archive is a gzipped tar file of an existing workspace. The archive contains all files in the workspace.

There are two ways to refer to an archive: 1) 2) //.

Workspace archives are stored in /var/ilx/workspaces//archive/. The system looks for archives in /var/ilx/workspaces//archive.

If an auth partition is removed from the BIG-IP configuration, archive files remain on disk and can be accessed by a user with the role of admin. A user with the role of admin can also manage archives under /var/ilx/workspaces in directories that do not correspond to a BIG-IP auth partition.

```
create workspace my_workspace archive my_workspace.tgz
```

Save the workspace as a gzipped tar file. The workspace must have been previously created.

```
create workspace my_workspace archive /my_partition/my_workspace.tgz
```

Save the workspace as a gzipped tar file to the set of archives in /my_partition. The workspace must have been previously created.

```
list workspace my_workspace archive
```

List all archives in the current auth partition.

```
list workspace my_workspace archive my_workspace.tgz
```

List file details for the archive my_workspace.tgz.

```
list workspace my_workspace archive /my_partition/my_workspace.tgz
```

List file details for the archive my_workspace.tgz that is in my_partition.

```
list workspace my_workspace archive /my_partition
```

List all archives in /my_partition.

```
delete workspace archive my_archive.tgz
```

Delete the workspace archive my_archive.tgz.

```
delete workspace archive /my_partition/my_archive.tgz
```

Delete the workspace archive my_archive.tgz in /my_partition.

```
create workspace my_workspace from-archive my_workspace.tgz
```

Create a new workspace from the archive my_workspace.tgz.

modify workspace my_workspace from-archive my_workspace.tgz

Update my_workspace from the archive my_workspace.tgz.

CONFIGURATION FILES

Workspaces are not stored in system config files.

Workspaces are stored in UCS files and are rolled forward to the next product release.

In a bladed system workspaces are clustered among blades.

Workspaces are not synced across device groups.

SEE ALSO

auth partition, cm device-group, create, delete, edit, ilx plugin, list, modify, pwd, sys folder, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2017-05-09 ilx workspace(1)

Itm

Itm alg-log-profile

NAME

alg-log-profile - Configures an Application-Level Gateway logging profile.

MODULE

itm

SYNTAX

CREATE/MODIFY

create alg-log-profile [name]

modify alg-log-profile [name | all]

options:

app-service [[string] | none]

csv-format [disabled | enabled]

start-control-channel {

action [disabled | enabled | backup-allocation-only]

elements [add | delete | replace-all-with] {

destination

}

}

end-control-channel {

action [disabled | enabled | backup-allocation-only]

elements [add | delete | replace-all-with] {

destination

}

}

start-data-channel {

action [disabled | enabled | backup-allocation-only]

elements [add | delete | replace-all-with] {

destination

}

}

end-data-channel {

action [disabled | enabled | backup-allocation-only]

elements [add | delete | replace-all-with] {

destination

}

}

inbound-transaction {

action [disabled | enabled]

}

edit alg-log-profile [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats alg-log-profile

reset-stats alg-log-profile [[name] | [glob] | [regex]] ...]

DISPLAY

```
list alg-log-profile
list alg-log-profile [ [name] | [glob] | [regex] ] ... ]
show running-config alg-log-profile
show running-config alg-log-profile [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

DELETE

```
delete alg-log-profile [name | all]
```

DESCRIPTION

A ALG log profile allows fine grain control of the logging for ALG events. When attached to an supported ALG profile - currently FTP, RTSP, SIP, you can control the events to log as well as optional elements in the log entry.

EXAMPLES

```
create ltm alg-log-profile my_alg_log_profile end-control-channel { action backup-allocation-only } end-data-channel { action backup-allocation-only }
```

Creates the ALG log profile `my_alg_log_profile` that generates log entries for both inbound and data-channel when translation is from backup members only.

```
delete alg-log-profile my_lsn_log_profile
```

Deletes the ALG log profile named `my_lsn_log_profile`.

OPTIONS

app-service

Specifies the name of the application service to which this object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

csv-format

When enabled, use CSV log format for log entries. The default value is disabled.

events

The type of ALG events available for logging control.

start-control-channel

Event for start of control channel connection.

end-control-channel

Event for end of control channel connection.

start-data-channel

Event for start of data channel connection.

end-data-channel

Event for end of data channel connection.

inbound-transaction

Event for inbound transaction event to an ALG end-point. Inbound transaction log entry could contain both incoming and outgoing messages.

action

Specify the logging action to be taken when a particular event is encountered.

enabled

Logging is enable for the event, regardless of how the flow is created.

backup-allocation-only

Logging is enable for the event, when the ALG is proxy with a LSN, and translation is take from backup pool member only.

disabled

Logging is disable for the event.

elements

Optional elements that can be added to the log message.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

ltm profile ftp, ltm profile sip, ltm profile rtsp, create, delete, edit, glob, list, ltm, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

ltm auth crldp-server

NAME

crldp-server - Creates a Certificate Revocation List Distribution Point (CRDLP) server for implementing a CRLDP authentication module.

MODULE

ltm auth

SYNTAX

Configure the crldp-server component within the ltm auth module using the syntax in the following sections.

CREATE/MODIFY

create crldp-server [name]

modify crldp-server [name]

options:

app-service [[string] | none]

base-dn [[LDAP base directory name] | none]

description [string]

host [[ip address] | none]

port [[name] | [number]]

reverse-dn [disabled | enabled]

edit crldp-server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list crldp-server

list crldp-server [[[name] | [glob] | [regex]] ...]

show running-config crldp-server

show running-config crldp-server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete crldp-server [name]

DESCRIPTION

CRLDP authentication is a mechanism for checking certificate revocation status for client connections passing through the BIG-IP(r) system. This module is useful when your authentication data is stored on a remote CRLDP server.

To implement a CRLDP authentication module and create a CRLDP server:

1. Use the crldp-server component in the ltm auth module to create a CRLDP server.
2. Use the ssl-crldp component in the ltm auth module to configure a CRLDP configuration object and associate it with the server you created in Step 1.
3. Use the profile component in the ltm auth module to create an authentication profile in which you specify the following options:
 - a. For the configuration option, specify the SSL CRLDP configuration object that you created in Step 2.
 - b. For the defaults-from option, specify a parent profile (either the default profile named ssl_crldp or another custom profile that you created).

EXAMPLES

```
create crldp-server my_crldp_server
```

Creates a CRLDP server named my_crldp_server.

```
delete crldp-server my_crldp_server
```

Deletes a CRLDP server named my_crldp_server.

OPTIONS

app-service

Specifies the name of the application service to which the CRLDP server belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the CRLDP server. Only the application service can modify or delete the CRLDP server.

base-dn

Specifies the LDAP base directory name for certificates that specify the CRL distribution point in directory name format (dirName). The default value is none.

Use this option when the value of the X509v3 attribute crlDistributionPoints is of type dirName. In this case, the BIG-IP system attempts to match the value of the crlDistributionPoints attribute to the value of the base-dn option. An example of a base-dn value is cn=lxxx,dc=f5,dc=com.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

host Specifies an IP address for the CRLDP server. This option is required. The default value is none.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

port Specifies the port for CRLDP authentication traffic. The default value is 389.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reverse-dn
Specifies in which order the system attempts to match the value of the base-dn option to the value of the X509v3 attribute crlDistributionPoints. When enabled, the system matches the value of the base-dn option from left to right, or from the beginning of the DN string, to accommodate dirName strings in certificates such as C=US,ST=WA,L=SEA,OU=F5,CN=xxx. The default value is disabled.

SEE ALSO

create, delete, edit, glob, list, ltm auth profile, ltm auth ssl-crl dp, ltm virtual, modify, reset-stats, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 ltm auth crldp-server(1)

ltm auth kerberos-delegation

NAME
kerberos-delegation - Configures a Kerberos delegation profile.

MODULE
ltm auth

SYNTAX
Configure the kerberos-delegation component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY
create kerberos-delegation [name]
modify kerberos-delegation [name]
options:
app-service [[string] | none]
client-principal [string]
debug-logging [disabled | enabled]
description [string]
protocol-transition [disabled | enabled]
server-principal [string]

edit kerberos-delegation [[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

reset-stats kerberos-delegation
reset-stats kerberos-delegation
[[name] | [glob] | [regex]] ...]

DISPLAY

```
list kerberos-delegation
list kerberos-delegation [ [ [name] | [glob] | [regex] ] ... ]
show running-config kerberos-delegation
show running-config kerberos-delegation
[ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

show kerberos-delegation
show kerberos-delegation [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global
```

DELETE

```
delete kerberos-delegation [name]
```

DESCRIPTION

The Kerberos delegation configuration acts like a proxy for Kerberos credentials. When connecting to a server that is inside its domain, the browser client fetches Kerberos credentials known as delegated credentials. These credentials are passed on to the system. Once the system has these credentials, it retrieves credentials for the RealServer(r) that is on the back end, and passes those credentials back.

Each user is assigned a unique cookie that describes a session on the system. This cookie is encrypted in a cookie key.

To configure a Kerberos authentication module and create a Kerberos configuration object:

1. Use the kerberos-delegation component in the ltm auth module to create a Kerberos configuration object.
2. Use the profile component, in the ltm auth module, to create an authentication profile in which you specify the following options:
 - a. For the configuration option, specify the Kerberos configuration object that you created in Step 1.
 - b. For the defaults-from option, specify a parent profile (either the default Kerberos profile named krbdelegate or another custom Kerberos profile that you created).

EXAMPLES

```
create kerberos-delegation my_kerberos-delegation_config client-principal client.net server-principal
server.net
```

Creates a Kerberos delegation profile named my_kerberos-delegation_config.

```
list kerberos-delegation all-properties
```

Displays all properties for all Kerberos delegation profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

client-principal

Specifies the principal that the client sees. This is usually a value such as HTTP/. This principal may be in a different domain from the server principal. This option is required. There is no default value.

debug-logging

Specifies whether the system logs debugging actions. The default value is disabled.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which this profile resides.

protocol-transition

Specifies whether associated virtual should transition client certificate authentication into Kerberos credentials.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

server-principal

Specifies the principal of the back-end web server. This is usually a value such as HTTP/. This may be in a different domain from the server

no default value.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-05-22 ltm auth kerberos-delegation(1)

ltm auth ldap

NAME

ldap - Configures an LDAP configuration object for implementing remote LDAP-based client authentication.

MODULE

ltm auth

SYNTAX

Configure the ldap component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY

create ldap [name]

modify ldap [name]

options:

bind-dn [[account dn] | none]

bind-pw [[string] | none]

bind-timeout [integer]

check-host-attr [disabled | enabled]

debug [disabled | enabled]

description [string]

filter [[string] | none]

group-dn [[group dn] | none]

group-member-attr [[string] | none]

idle-timeout [integer]

ignore-auth-info-unavail [no | yes]

ignore-unknown-user [disabled | enabled]

login-attribute [[account name] | none]

port [[name] | [integer]]

scope [base | one | sub]

search-base-dn [[search base dn] | none]

search-timeout [number]

servers

[add | delete | replace-all-with] {

[ip address ...]

}

servers none

ssl [disabled | enabled]

ssl-ca-cert-file [[name] | none]

ssl-check-peer [disabled | enabled]

ssl-ciphers [[string] | none]

ssl-client-cert [[string] | none]

ssl-client-key [[string] | none]

user-template [[string] | none]

version [number]

warnings [disabled | enabled]

edit ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ldap

list ldap [[[name] | [glob] | [regex]] ...]

show running-config ldap

show running-config ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete ldap [name]

DESCRIPTION

LDAP authentication is a mechanism for authenticating or authorizing client connections passing through the system. LDAP authentication is useful when your authentication or authorization data is stored on a remote LDAP server or a Microsoft(r) Windows Active Directory(r) server, and you want the client credentials to be based on basic HTTP authentication (that is, user name and password).

To configure an LDAP authentication module and create an LDAP configuration object:

1. Use the ldap component in the ltm auth module to create an LDAP configuration object.
2. Use the profile component, in the ltm auth module, to create an authentication profile in which you specify the following options:
 - a. For the configuration option, specify the LDAP configuration object that you created in Step 1.
 - b. For the defaults-from option, specify a parent profile (either the default LDAP profile named ldap or another custom profile that you created).

EXAMPLES

```
create ldap my_auth_ldap servers add {my_ldap_auth_server}
```

Creates a configuration object named my_auth_ldap

```
delete ldap my_auth_ldap
```

Deletes the configuration object named my_auth_ldap.

OPTIONS

bind-dn

Specifies the distinguished name of an account to which to bind, to perform searches. This search account is a Read-only account used to do searches. You can use the admin account as the search account. If no admin DN is specified, then no bind is attempted. The default value is none.

This option is required only when a site does not allow anonymous searches. If the remote server is a Microsoft(r) Windows(r) Active Directory(r) server, the distinguished name must be in the form of an email address.

bind-pw

Specifies the password for the search account created on the LDAP server. This option is required if you specify a value for the bind-dn option. The default value is none.

bind-timeout

Specifies a bind timeout limit. The default value is 30 seconds.

check-host-attr

Confirms the password for the bind distinguished name. This option is optional. The default value is disabled.

debug

Enables or disables syslog-ng debugging information at LOG DEBUG level. The default value is disabled. F5 Networks does not recommend using this option for normal configuration.

description

User defined description.

filter

Specifies a filter. Use this option for authorizing client traffic. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

group-dn

Specifies the group distinguished name. The system uses this option for authorizing client traffic. The default value is none.

group-member-attribute

Specifies a group member attribute. The system uses this option for authorizing client traffic. The default value is none.

idle-timeout

Specifies the idle timeout, in seconds, for connections. The default value is 3600 seconds.

ignore-auth-info-unavail

Specifies whether the system ignores authentication information, if it is not available. The default value is no.

ignore-unknown-user

Specifies whether the system ignores a user that is unknown. The default value is disabled.

login-attribute

Specifies a logon attribute. Normally, the value for this option is uid; however, if the server is a Microsoft Windows Active Directory server, the value must be the account name samaccountname (not case-sensitive). The default value is none.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

`port` Specifies the port number or name for the LDAP service. Port 389 is typically used for non-SSL and port 636 is used for an SSL-enabled LDAP service. The default value is `ldap`.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`scope`

Specifies the search scope. The default value is `sub`. The options are:

`base` Specifies the search scope is base object. The base value is almost never useful for nameservice lookups.

`one` Specifies the search scope is one level.

`sub` Specifies the search scope is subtree.

`search-base-dn`

Specifies the search base distinguished name. The default value is `none`.

`search-timeout`

Specifies the search timeout. The default value is 30 seconds.

`servers`

Specifies the LDAP servers that the system must use to obtain authentication information. You must specify a server when you create an LDAP configuration object.

`ssl` Enables or disables SSL functionality. The default is disabled.

Note that when you use the command line interface to enable SSL for an LDAP service, the system does not change the service port number from 389 to 636, as is required. To change the port number from the command line, use the `service` option of this command (see above), for example, `ldap [name] ssl enabled service 636`.

`ssl-ca-cert-file`

Specifies the name of an SSL CA certificate using the full path to the file. The default value is `none`.

`ssl-check-peer`

Specifies whether the system checks an SSL peer. The default value is disabled.

`ssl-ciphers`

Specifies SSL ciphers. The default value is `none`.

`ssl-client-cert`

Specifies the name of an SSL client certificate. The default value is `none`.

`ssl-client-key`

Specifies the name of an SSL client key. The default value is `none`.

`user-template`

Specifies a user template for the LDAP application to use for authentication. The default value is `none`.

`version`

Specifies the version number of the LDAP application. The default value is 3.

`warnings`

Enables or disables warning messages. The default value is enabled.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm auth profile`, `ltm virtual`, `modify`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm auth ldap(1)

ltm auth ocsponder

NAME

`ocsponder` - Configures Online Certificate System Protocol (OCSP) responder objects.

MODULE

`ltm auth`

SYNTAX

Configure the ocsdp-responder component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ocsdp-responder [name]
modify ocsdp-responder [name]
options:
  allow-certs [disabled | enabled]
  app-service [[string] | none]
  ca-file [ [file name] | none]
  ca-path [ [file name] | none]
  cert-id-digest [md5 | sha1]
  chain [disabled | enabled]
  check-certs [disabled | enabled]
  description [string]
  explicit [disabled | enabled]
  ignore-aia [disabled | enabled]
  intern [disabled | enabled]
  nonce [disabled | enabled]
  sign-digest [md5 | sha1]
  sign-key [ [key] | none]
  sign-key-pass-phrase [ [pass phrase] | none]
  sign-other [ [list of certs] | none]
  signer [ [certificate] | none]
  status-age [integer]
  trust-other [disabled | enabled]
  url [none | [url] ]
  va-file [ [file name] | none]
  validity-period [integer]
  verify [disabled | enabled]
  verify-cert [disabled | enabled]
  verify-other [ [file name] | none]
  verify-sig [disabled | enabled]
```

```
edit ocsdp-responder [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list ocsdp-responder
list ocsdp-responder [ [ [name] | [glob] | [regex] ] ... ]
show running-config ocsdp-responder
show running-config ocsdp-responder [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete ocsdp-responder [name]
```

DESCRIPTION

To implement the SSL OCSF authentication module, you must create the following objects: one or more OCSF responder objects, an SSL OCSF configuration object, and an SSL OCSF profile.

To implement an SSL OCSF authentication module and create an OCSF responder object:

1. Use the ocsdp-responder component in the ltm auth module to configure an OCSF responder object.
2. Use the ssl-ocsf component in the ltm auth module to configure an SSL OCSF configuration object to which you add the OCSF responder object that you created in Step 1.
3. Use the profile component in the ltm auth module to create an authentication profile in which you specify the following options:
 - a. For the configuration option, specify the SSL OCSF configuration object that you created in Step 2.
 - b. For the defaults-from option, specify a parent profile (either the default OCSF Responder profile named `ssl_ocsf` or another custom profile that you created).

OPTIONS

allow-certs
Enables or disables the addition of certificates to an OCSF request. The default value is enabled.

app-service
Specifies the name of the application service to which the OCSF responder object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the OCSF responder object. Only the application service can modify or delete the OCSF responder object.

ca-file
Specifies the name of the file containing trusted CA certificates used to verify the signature on the OCSF response. The default value is none.

ca-path
Specifies the name of the path containing trusted CA certificates used to verify the signature on the OCSF response. The default value is none.

cert-id-digest

Specifies a specific algorithm identifier, either sha1 or md5. The default value is sha1. The options are:

sha1 is newer and provides more security with a 160-bit hash length.
md5 is older and has only a 128-bit hash length.

The cert ID is part of the OCSP protocol. The OCSP client (in this case, the BIG-IP system) calculates the cert ID using a hash of the Issuer and serial number for the certificate that it is trying to verify.

chain

Specifies whether the system constructs a chain from certificates in the OCSP response. The default value is enabled.

check-certs

Enables or disables verification of an OCSP response certificate. Use this option for debugging purposes only. The default value is enabled.

description

User defined description.

explicit

Specifies that the Local Traffic Manager explicitly trusts that the OCSP response signer's certificate is authorized for OCSP response signing. If the signer's certificate does not contain the OCSP signing extension, specification of this option causes a response to be untrusted. The default value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-aia

Specifies whether the system ignores the URL contained in the certificate's AIA fields, and always uses the URL specified by the responder instead. The default value is disabled.

intern

Specifies whether the system ignores certificates contained in an OCSP response when searching for the signer's certificate. To use this option, the signer's certificate must be specified with either the verify-other or va-file option. The default value is enabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

nonce

Specifies whether the system verifies an OCSP response signature or the nonce values. The default value is enabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

sign-digest

Specifies the algorithm for signing the request, using the signing certificate and key. This parameter has no meaning, if request signing is not in effect (that is, both the request signing certificate and request signing key parameters are empty). This parameter is required only when request signing is in effect. The default value is sha1.

sign-key

Specifies the key that the system uses to sign an OCSP request. The default value is none.

sign-key-pass-phrase

Specifies the passphrase that the system uses to encrypt the sign key. The default value is none.

sign-other

Adds a list of additional certificates to an OCSP request. The default value is none.

signer

Specifies a certificate used to sign an OCSP request. If the certificate is specified, but the key is not specified, then the private key is read from the same file as the certificate. If neither the certificate nor the key is specified, then the request is not signed. If the certificate is not specified and the key is specified, then the configuration is considered to be invalid. The default value is none.

status-age

Specifies the age of the status of the OCSP responder. The default value is 0 (zero).

trust-other

Instructs the BIG-IP local traffic management system to trust the BIG certificates specified with the verify-other option. The default is value disabled.

url Specifies the URL used to contact the OCSP service on the responder. This option is required. The default value is none.

va-file

Specifies the name of the file containing explicitly trusted responder certificates. This parameter is

needed in the event that the responder is not covered by the certificates already loaded into the responder's CA store. The default value is none.

validity-period

Specifies the number of seconds used to specify an acceptable error range. Use this option when the OCSP responder clock and a client clock are not synchronized, which can cause a certificate status check to fail. This value must be a positive number. The default value is 300 seconds.

verify

Enables or disables verification of an OCSP response signature or the nonce values. Used for debugging purposes only. The default value is enabled.

verify-cert

Specifies that the system makes additional checks to see if the signer's certificate is authorized to provide the necessary status information. Use this option for testing purposes only. The default value is enabled.

verify-other

Specifies the name of the file used to search for an OCSP response signing certificate when the certificate has been omitted from the response. The default value is none.

verify-sig

Specifies that the system checks the signature on the OCSP response. Use this option for testing purposes only. The default value is enabled.

SEE ALSO

create, delete, edit, glob, list, ltm auth profile, ltm auth ssl-ocsp, ltm virtual, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2012-10-19 ltm auth ocsp-responder(1)

ltm auth profile

NAME

profile - Configures an authentication profile.

MODULE

ltm auth

SYNTAX

Configure the profile component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY

create profile [name]

modify profile [name]

options:

app-service [[string] | none]

configuration [[name] | none]

cookie-key [string]

cookie-name [string]

credential-source [http-basic-auth]

defaults-from [name]

description [string]

enabled [yes | no]

idle-timeout [integer]

rule [iRule name]

edit profile [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats profile

reset-stats profile [[[name] | [glob] | [regex]] ...]

DISPLAY

list profile

list profile [[[name] | [glob] | [regex]] ...]

show running-config profile

show running-config profile [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties
one-line
partition

show profile
show profile [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete profile [name]

Note: You cannot delete default profiles.

DESCRIPTION

You can use the profile component to configure a custom authentication profile, or you can use the default profile that the BIG-IP(r) Local Traffic Manager system provides for each type of authentication module.

An authentication profile requires one of the following configuration objects: ltm auth kerberos-delegation, ltm auth ldap, ltm auth radius, ltm auth ssl-cc-ldap, ltm auth ssl-crldp, ltm auth ssl-ocsp or ltm auth tacacs. The type of profile specified by the defaults-from option must match the type of configuration object.

EXAMPLES

```
create profile my_authentication_profile { configuration tacacs defaults-from tacacs credential-source http-basic-auth enabled yes idle-timeout 30 rule _sys_auth_tacacs }
```

Creates a profile named my_authentication_profile for TACACS+ authentication.

```
list profile
```

Displays the properties of all of the auth profile components.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

configuration

Specifies the name of an authentication configuration object. This option is required.

cookie-key

Specifies the key that the system uses to encrypt the session cookie assigned to each user using the cookie-name option. The default value is f5auth. This option applies only to KRB Delegate profiles.

cookie-name

Specifies a session cookie that the system assigns to each user. F5 Networks recommends that each virtual server use a different cookie name. The system encrypts the cookie using the value of the cookie-key option. The default value is abc123. This option applies only to KRB Delegate profiles.

credential-source

Specifies the credential source.

defaults-from

Specifies the name of the authentication profile from which you want your custom profile to inherit settings. This option is required.

description

User defined description.

enabled

Specifies whether this authentication profile is enabled. The default value is yes.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

idle-timeout

Specifies the idle timeout for the authentication profile. The default value is 300 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rule Specifies the name of the rule that corresponds to the authentication method you want to use.

SEE ALSO

create, delete, edit, glob, ltm auth crldp-server, ltm auth kerberos-delegation, ltm auth ldap, ltm auth ocsp-responder, ltm auth radius, ltm auth radius-server, ltm auth ssl-cc-ldap, ltm auth ssl-crldp, ltm auth ssl-

ocsp, ltm auth tacacs, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm auth profile(1)

ltm auth radius-server

NAME

radius-server - Configures a RADIUS server for implementing remote RADIUS-based client authentication.

MODULE

ltm auth

SYNTAX

Configure the radius-server component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY

create radius-server [name]

modify radius-server [name]

options:

description [string]

port [[name] | [number]]

secret [none | ["string"]]

server [[hostname] | [ip address] | none]

timeout [integer]

edit radius-server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list radius-server

list radius-server [[[name] | [glob] | [regex]] ...]

show running-config radius-server

show running-config radius-server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete radius-server [name]

DESCRIPTION

You use a RADIUS authentication module when your authentication data is stored on a remote RADIUS server. In this case, client credentials are based on basic HTTP authentication (that is, user name and password).

To configure a RADIUS authentication module and create a RADIUS server:

1. Use the radius-server component in the ltm auth module to configure a RADIUS server.
2. Use the radius component in the ltm auth module to create a RADIUS configuration object that references the RADIUS server you created in Step 1.
3. Use the profile component in the ltm auth module to create an authentication profile in which you specify the following options:
 - a. For the configuration option, specify the radius component that you created in Step 2.
 - b. For the defaults-from option, specify a parent profile (either the default RADIUS profile named radius or another custom profile that you created).

EXAMPLES

```
create radius-server bigip_auth_radius_server secret "This is the secret." server 10.1.1.1
```

Creates a RADIUS server named my_radius_server.

```
delete radius-server my_radius_server
```

Deletes the RADIUS server named my_radius_server.

OPTIONS

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition in which the component resides.

port Specifies the port for RADIUS authentication traffic. The default value is 1812.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret
Specifies the secret key the system uses to encrypt and decrypt packets sent or received from the server. This option is required.

server
Specifies the host name or IP address of the RADIUS server. This option is required.

timeout
Specifies the timeout value. The default value is 3 seconds.

SEE ALSO

create, delete, edit, glob, list, ltm auth profile, ltm auth radius, ltm virtual, modify, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-04-06 ltm auth radius-server(1)

ltm auth radius

NAME

radius - Configures a RADIUS configuration object for implementing remote RADIUS-based authentication of BIG-IP(r) system users.

MODULE

ltm auth

SYNTAX

Configure the radius component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY

create radius [name]

modify radius [name]

options:

accounting-bug [disabled | enabled]

client-id [none | [string]]

debug [disabled | enabled]

description [string]

retries [integer]

service-type [default | login | framed | callback-login | callback-framed | outbound | administrative | nas-prompt | authenticate-only | callba

servers

[add | delete | replace-all-with] {

[[hostname ...] | [ip address ...]]

}

servers [default | none]

edit radius [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list radius

list radius [[[name] | [glob] | [regex]] ...]

show running-config radius

show running-config radius [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

DELETE

delete radius [name]

DESCRIPTION

You use a RADIUS authentication module when your authentication data is stored on a remote RADIUS server. In this case, client credentials are based on basic HTTP authentication (that is, username and password).

To implement a RADIUS authentication module and create a RADIUS configuration object:

1. Use the radius-server component in the ltm auth module to configure a RADIUS server.
2. Use the radius component in the ltm auth module to create a RADIUS configuration object that references the RADIUS server you created in Step 1.
3. Use the profile component in the ltm auth module to create an authentication profile in which you specify the following options:
 - a. For the configuration option, specify the RADIUS configuration object that you created in Step 2.
 - b. For the defaults-from option, specify a parent profile (either the default RADIUS profile named radius or another custom profile that you created).

EXAMPLES

```
create radius my_radius_auth servers add { myradiusserver }
```

Creates a RADIUS configuration object named my_radius_auth.

```
delete radius my_radius_auth
```

Deletes the RADIUS configuration object named my_radius_auth.

OPTIONS

accounting-bug

Enables or disables validation of the accounting response vector. This option is necessary only on older servers. The default value is disabled.

client-id

Sends a NAS-Identifier RADIUS attribute with string bar. If you do not specify a value for the client-id option, the system uses the pluggable authentication module (PAM) service type. You can disable this feature by specifying a blank client ID.

debug

Enables or disables syslog-ng debugging information at LOG DEBUG level. F5 Networks does not recommend this option for normal use. The default value is disabled.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

retries

Specifies the number of authentication retries that the Local Traffic Manager allows before authentication fails. The default value is 3.

service-type

Specifies the type of service used for the RADIUS server. The default is default, which behaves as authenticate-only.

servers

Specifies the hostnames or IP addresses of the RADIUS servers that the BIG-IP Local Traffic Manager uses to obtain authentication data.

SEE ALSO

create, delete, edit, glob, list, ltm auth profile, ltm auth radius-server, ltm virtual, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

Itm auth ssl-cc-ldap

NAME

ssl-cc-ldap - Configures an SSL client certificate configuration object for remote SSL-based LDAP authorization for client traffic passing through the traffic management system.

MODULE

itm auth

SYNTAX

Configure the ssl-cc-ldap component within the Itm auth module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ssl-cc-ldap [name]
modify ssl-cc-ldap [name]
options:
  admin-dn [ [name] | none]
  admin-password [none | [password] ]
  cache-size [integer]
  cache-timeout [integer]
  certmap-base [none | [search base] ]
  certmap-key [ [name] | none]
  certmap-user-serial [no | yes]
  description [string]
  group-base [none | [search base] ]
  group-key [ [name] | none]
  group-member-key [[name] | none]
  role-key [ [name] | none]
  search-type [cert | certmap | user]
  secure [no | yes]
  servers
    [add | delete | none | replace-all-with] {
[ip address ... ]
}
  user-base [none | [search base] ]
  user-class [ [class] | none]
  user-key [ [key] | none]
  valid-groups
    [add | delete | replace-all-with] {
[group ... ]
}
  valid-groups none
  valid-roles
    [add | delete | replace-all-with] {
[role ... ]
}
  valid-roles none
```

```
edit ssl-cc-ldap [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list ssl-cc-ldap
list ssl-cc-ldap [ [ [name] | [glob] | [regex] ] ... ]
show running-config ssl-cc-ldap
show running-config ssl-cc-ldap
[ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete ssl-cc-ldap [name]
```

DESCRIPTION

You can use the ssl-cc-ldap component to configure SSL client certificate-based remote LDAP authorization for client traffic passing through the traffic management system.

To configure this type of authentication module and create a configuration object:

1. Use the ssl-cc-ldap component in the Itm auth module to create an SSL client certificate LDAP configuration object.
2. Use the profile component in the Itm auth module to create an authentication profile in which you specify the following options:

- a. For the configuration option, specify the configuration object that you created in Step 1.
- b. For the defaults-from option, specify a parent profile (either the default profile named `ssl_cc_ldap` or another custom profile that you created).

OPTIONS

`admin-dn`

Specifies the distinguished name of an account to which to bind to perform searches. This search account is a read-only account used to do searches. The admin account can also be used as the search account. If no admin DN is specified, then no bind is attempted.

This option is required only when an LDAP database does not allow anonymous searches. The default value is none.

`admin-password`

Specifies the password for the admin account. See `admin-dn` above. The default value is none.

`cache-size`

Specifies the maximum size, in bytes, allowed for the SSL session cache. Setting this option to 0 (zero) disallows SSL session caching. The default value is 20000 bytes (20KB).

`cache-timeout`

Specifies the number of usable lifetime seconds of negotiable SSL session IDs. When this time expires, a client must negotiate a new session. The default value is 300 seconds.

`certmap-base`

Specifies the search base for the subtree used by the certmap search method. A typical search base is: `ou=people,dc=company,dc=com`. The default value is none.

`certmap-key`

Specifies the name of the certificate map that the certmap search method uses. This name is found in the LDAP database. The default value is none.

`certmap-user-serial`

Specifies whether the system uses the client certificate's subject or serial number (in conjunction with the certificate's issuer) when trying to match an entry in the certificate map subtree.

A value of `yes` uses the serial number. A value of `no` uses the subject. The default value is `no`.

`description`

User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`group-base`

Specifies the search base for the subtree used by group searches. Use this option only when specifying the `valid-groups` option. The typical search base is similar to: `ou=groups,dc=company,dc=com`. The default value is none.

`group-key`

Specifies the name of the attribute in the LDAP database that specifies the group name in the group subtree. An example of a typical key is `cn` (common name for the group). The default value is none.

`group-member-key`

Specifies the name of the attribute in the LDAP database that specifies members (DNs) of a group. A typical key is `member`. The default value is none.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`role-key`

Specifies the name of the attribute in the LDAP database that specifies a user's authorization roles. Use this option only when specifying the `valid-roles` option. A typical role key is `authorizationRole`. The default value is none.

`search`

Specifies the type of LDAP search that is performed based on the client's certificate. Possible values are:

`cert` Searches for the exact certificate.

`certmap`

Searches for a user by matching the certificate issuer and the certificate serial number or certificate.

`user` Searches for a user based on the common name found in the certificate. This is the default value.

`secure`

Specifies whether the system attempts to use secure LDAP (LDAP over SSL). The alternative to using secure

LDAP is to use insecure (clear text) LDAP. Secure LDAP is a consideration when the connection between the BIG-IP system and the LDAP server cannot be trusted. The default value is no.

servers

Specifies a list of LDAP servers you want to search. You must specify a server when you create an SSL client certificate configuration object.

user-base

Specifies the search base for the subtree used when you select for the search option either of the values user or cert. A typical search base is: ou=people,dc=company,dc=com. You must specify a user base when you create an SSL client certificate configuration object. The default value is none.

user-class

Specifies the object class in the LDAP database to which the user must belong to be authenticated. The default value is none.

user-key

Specifies the key that denotes a user ID in the LDAP database (for example, the common key for the user option is uid). You must specify a user key when you create an SSL client certificate configuration object.

valid-groups

Specifies a space-delimited list of the names of groups to which the client must belong in order to be authorized (matches against the group key in the group subtree). The client needs to be a member of only one of the groups in the list. The default value is none.

valid-roles

Specifies a space-delimited list of the valid roles that clients must have to be authorized. The default value is none.

SEE ALSO

create, delete, edit, glob, list, ltm auth profile, ltm virtual, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 ltm auth ssl-cc-ldap(1)

ltm auth ssl-crl dp

NAME

ssl-crl dp - Configures a Secure Socket Layer (SSL) Certificate Revocation List Distribution Point (CRLDP) configuration object for implementing SSL CRLDP to manage certificate revocation.

MODULE

ltm auth

SYNTAX

Configure the ssl-crl dp component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ssl-crl dp [name]
```

```
modify ssl-crl dp [name]
```

options:

```
cache-timeout [integer]
```

```
connection-timeout [integer]
```

```
description [string]
```

servers

```
[add | delete | replace-all-with] {  
[ip address ... ]  
}
```

```
servers [default | none]
```

```
update-interval [integer]
```

```
use-issuer [disabled | enabled]
```

```
edit ssl-crl dp [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list ssl-crl dp
```

```
list ssl-crl dp [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config ssl-crl dp
```

```
show running-config ssl-crl dp
```

[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line
- partition

DELETE

delete ssl-crl dp [name]

DESCRIPTION

CRLDP authentication is a mechanism for checking certificate revocation status for client connections passing through the system. This module is useful when your authentication data is stored on a remote CRLDP server.

To implement a CRLDP authentication module and create an SSL CRLDP configuration object:

1. Use the `crl dp-server` component, in the `l tm auth` module, to create a CRLDP server.
2. Use the `ssl-crl dp` component in the `l tm auth` module to configure a CRLDP configuration object that references the server you created in Step 1.
3. Use the `profile` component in the `l tm auth` module to create an authentication profile in which you specify the following options:
 - a. For the `configuration` option, specify the SSL CRLDP configuration object that you created in Step 2.
 - b. For the `defaults-from` option, specify a parent profile (either the default profile named `ssl_crl dp` or another custom profile that you created).

EXAMPLES

```
create ssl-crl dp my_auth_ssl-crl dp
```

Creates an SSL CRLDP configuration object named `my_auth_ssl-crl dp`.

```
delete ssl-crl dp my_auth_ssl-crl dp
```

Deletes the SSL CRLDP configuration object named `my_auth_ssl-crl dp`.

OPTIONS

`cache-timeout`

Specifies the number of seconds that CRLs are cached. The default value is 86400 (24 hours).

`connection-timeout`

Specifies the number of seconds before the connection times out. The default value is 15.

`description`

User defined description.

`glob` Displays the items that match the `glob` expression. See help `glob` for a description of `glob` expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an `@` sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`servers`

Specifies a host name or IP address for the secure CRLDP server. This option is required. The default value is none.

`update-interval`

Specifies an update interval for CRL distribution points that ensures that CRL status is checked at regular intervals, regardless of the CRL timeout value. This helps to prevent CRL information from becoming outdated before the BIG-IP system checks the status of a certificate. The default value is 0 (zero), which indicates an internal default value is active.

`use-issuer`

Specifies whether the system extracts the CRL distribution point from the client certificate. The default value is disabled.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `l tm auth profile`, `l tm auth crl dp-server`, `l tm virtual`, `modify`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

Itm auth ssl-ocsp

NAME

ssl-ocsp - Configures OCSP authentication for client traffic passing through the traffic management system.

MODULE

itm auth

SYNTAX

Configure the ssl-ocsp component within the Itm auth module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ssl-ocsp [name]
modify ssl-ocsp [name]
options:
  description [string]
  responders
    [add | delete | replace-all-with] {
[name]...
}
  responders [default | none]
```

```
edit ssl-ocsp [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list ssl-ocsp
list ssl-ocsp [ [ [name] | [glob] | [regex] ] ... ]
show running-config ssl-ocsp
show running-config ssl-ocsp
[ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties
one-line
partition
```

DELETE

```
delete ssl-ocsp [name]
```

DESCRIPTION

Online Certificate Status Protocol (OCSP) is an industry-standard protocol that offers an alternative to a certificate revocation list when using public-key technology. To implement an SSL OCSP authentication module, you must create the following objects: one or more OCSP responder objects, an SSL OCSP configuration object, and an SSL OCSP profile.

To implement an SSL OCSP authentication module and create an SSL OCSP configuration object:

1. Use the obsp-responder component in the Itm auth module to configure an OCSP responder object.
2. Use the ssl-ocsp component in the Itm auth module to configure an SSL OCSP configuration object to which you add the OCSP responder object that you created in Step 1.
3. Use the profile component in the Itm auth module to create an authentication profile in which you specify the following options:
 - a. For the configuration option, specify the SSL OCSP configuration object that you created in Step 2.
 - b. For the defaults-from option, specify a parent profile (either the default OCSP Responder profile named `ssl_ocsp` or another custom profile that you created).

EXAMPLES

```
create ssl-ocsp my_auth_ssl-ocsp
```

Creates an SSL OCSP configuration object named `my_auth_ssl-ocsp`.

```
delete ssl-ocsp my_auth_ssl-ocsp
```

Deletes the SSL OCSP configuration object named `my_auth_ssl-ocsp`.

OPTIONS

`description`
User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`
Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

responders
Specifies a list of OCSP responders that you configured using the ocspp-responder component in the ltm auth module.

SEE ALSO

create, delete, edit, glob, list, ltm auth profile, ltm auth ocspp-responder, ltm virtual, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 ltm auth ssl-ocsp(1)

ltm auth tacacs

NAME

tacacs - Configures a TACACS+ configuration component for implementing remote TACACS+-based client authentication.

MODULE

ltm auth

SYNTAX

Configure the tacacs component within the ltm auth module using the syntax shown in the following sections.

CREATE/MODIFY

create tacacs [name]

modify tacacs [name]

options:

accounting [send-to-all-servers | send-to-first-server]

authentication [use-all-servers | use-first-server]

debug [disabled | enabled]

description [string]

encryption [disabled | enabled]

protocol [none | [protocol]]

secret ["[string]"]

servers

[add | delete | replace-all-with] {

[[[hostname[:port]] | [ip address[:port]]] ...]

}

service [[name] | none]

edit tacacs [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tacacs

list tacacs [[[name] | [glob] | [regex]] ...]

show running-config tacacs

show running-config tacacs [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete tacacs [name]

DESCRIPTION

Using a TACACS+ configuration object and profile, you can implement the TACACS+ authentication module as the mechanism for authenticating client connections passing through the BIG-IP Local Traffic Manager system. You use this module when your authentication data is stored on a remote TACACS+ server. In this case, client credentials are based on basic HTTP authentication (that is, user name and password).

To implement a TACACS+ authentication module and create a TACACS configuration object:

1. Use the tacacs component in the ltm auth module to configure a TACACS+ configuration object.
2. Use the profile component in the ltm auth module to create an authentication profile in which you specify

the following options:

- a. For the configuration option, specify the TACACS+ configuration object that you created in Step 1.
- b. For the defaults-from option, specify a parent profile (either the default TACACS+ profile named tacacs or another custom profile that you created).

EXAMPLES

```
create tacacs my_tacacs_auth secret "This is the secret" servers add {my_tacacs_server} encryption enabled
```

Enables encryption for TACACS+ packets.

```
create tacacs my_tacacs_auth secret "This is the secret" servers add { my_tacacs_server1 my_tacacs_server2 }  
accounting send-to-all-servers
```

Provides the ability to send accounting start and stop packets to all servers

OPTIONS

accounting

If multiple TACACS+ servers are defined and pluggable authentication module (PAM) session accounting is available, specifies where the system sends accounting start and stop packets. Possible values are:

send-to-all-servers

Sends to all servers.

send-to-first-server

Sends to the first available server.

authentication

Specifies when to use the secret key supplied for the secret option. This option is required. The options are:

use-all-servers

Use the secret key with all servers.

use-first-server

Use the secret key with the first available server.

debug

Enables syslog-ng debugging information at LOG DEBUG level. F5 Networks does not recommend this option for normal use. The default value is disabled.

description

User defined description.

encryption

Enables or disables encryption of TACACS+ packets. F5 Networks recommends this option for normal use. The default value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

protocol

Specifies the protocol associated with the value specified in the service option, which is a subset of the associated service being used for client authorization or system accounting.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Sets the secret key used to encrypt and decrypt packets sent or received from the server. This option is required.

servers

Specifies the host name or IPv4 address of the TACACS+ server. For each server, a port may optionally be specified in the format hostname:port or IPv4:port. If no port is specified, the default port 49 is used. This option is required.

service

Specifies the name of the service that the user is requesting to be authenticated to use. Identifying the service enables the TACACS+ server to behave differently for different types of authentication requests. This option is required.

SEE ALSO

create, delete, edit, glob, list, ltm auth profile, ltm virtual, modify, regex, show, tmsh,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

Itm cipher group

NAME

group - Configures a cipher group.

MODULE

itm cipher

SYNTAX

Configure the group component within the cipher module using the syntax shown in the following sections.

CREATE/MODIFY

```
create group [name]
modify group [name]
options:
  allow { add rule1 [ rule2... ] }
  description [string]
  exclude { add rule3 [ rule4... ] }
  require { add rule5 [ rule6... ] }
```

```
edit group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list group
list group [ [ [name] | [glob] | [regex] ] ... ]
```

```
show group [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete group [all | [name]]
options:
  recursive
```

DESCRIPTION

You can use the group component to create, modify, or delete a custom cipher group, or display a custom cipher group.

Cipher groups are contain sets of cipher rules and are attached to client-ssl or server-ssl profiles.

EXAMPLES

```
create group my_group { allow add { f5-default } }
```

Creates a group named my_group with a single allowed rule, f5-default.

OPTIONS

allow
Specifies a list of rules that are allowed in this group.

description
User defined description.

exclude
Specifies a list of rules that are excluded from this group.

require
Specifies a list of rules that are required for this group. The operation is the intersection of allow and require after all excluded rules have been removed.

SEE ALSO

create, delete, edit, glob, list, Itm virtual, modify, mv, regex, reset-stats, show, tmsd

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

Itm cipher rule

NAME

rule - Configures a cipher rule.

MODULE

itm cipher

SYNTAX

Configure the rule component within the cipher module using the syntax shown in the following sections.

CREATE/MODIFY

create rule [name]

modify rule [name]

options:

cipher [string]

description [string]

dh-groups [string]

signature-algorithms [string]

edit rule [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list rule

list rule [[name] | [glob] | [regex]] ...]

show rule [[name] | [glob] | [regex]] ...]

DELETE

delete rule [all | [name]]

options:

recursive

DESCRIPTION

You can use the rule component to create, modify, or delete a custom cipher rule, or display a custom cipher rule.

Cipher rules are gathered into cipher groups and attached to client-ssl or server-ssl profiles.

EXAMPLES

```
create rule my_rule cipher "default"
```

Creates a rule named my_rule with a cipher string "default".

OPTIONS

cipher rule

Specifies the OpenSSL compatible cipher string.

description

User defined description.

dh-groups groups

Specifies the allowed named groups, separated by ":". For example: "P256:X25519"

The available named groups are: P256, P384, X25519 A special keyword, DEFAULT, represents the recommended set of named groups.

signature-algorithms signature algorithms

Specifies the allowed signatures algorithms, separated by ":". For example:

"RSA_PKCS1_SHA256:ECDSA_P256_SHA256"

The available signature algorithms are: DSA-SHA1, DSA-SHA256, DSA-SHA384, DSA-SHA512, ECDSA-SHA1, ECDSA-SHA256, ECDSA-SHA384, ECDSA-SHA512, RSA-PKCS1-SHA1, RSA-PKCS1-SHA256, RSA-PKCS1-SHA384, RSA-PKCS1-SHA512, RSA-PSS-SHA256, RSA-PSS-SHA384, RSA-PSS-SHA512

A special keyword, DEFAULT, represents the recommended set of signature algorithms.

SEE ALSO

create, delete, edit, glob, list, Itm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

Itm classification application

NAME

application - Configures a custom classification application.

MODULE

itm classification

SYNTAX

Configure the application within the itm classification module using the syntax shown in the following sections.

CREATE/MODIFY

create application [name]

options:

app-service [[string] | none]

description [string]

application-id [integer]

category [name]

modify application [name]

options:

app-service [[string] | none]

description [string]

category [name]

edit application [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list application

list application [[all] | [name]]

show running-config application

show running-config application [[all] | [name]]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete application [name]

Note: All referring classification-filters (to this application) need to be deleted first; otherwise an error will be reported. Predefined applications cannot be deleted.

DESCRIPTION

You can use the application component to create, modify, delete, and display classification application.

EXAMPLES

```
create application my_app { application-id 8192 category my_cat }
```

Creates a new application named my_app.

```
modify application my_app { category Web description "My description." }
```

Modify an application named my_app.

```
list application
```

Displays all created applications.

```
delete application my_app
```

Deletes the application named my_app.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

application-id

Identifies the application. This is set during creation and cannot be changed. Identifiers must be unique

across predefined and user-defined applications. Predefined application-ids must be in numeric range [0, 8192), and user defined application-ids must be in numeric range [8192, 16384).

category

Refers to classification category. The referred category [name] should exist already; otherwise an error will be reported.

SEE ALSO

create, modify, delete, list, show, tmsh, ltm classification

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm classification application(1)

ltm classification auto-update settings

NAME

settings - Configures settings for classification hitless upgrade.

MODULE

ltm classification classification auto-update settings

SYNTAX

MODIFY

modify ltm classification auto-update settings

options:

auto-update-interval [daily|monthly|weekly]

enabled [yes|no]

DISPLAY

list ltm classification auto-update settings

DESCRIPTION

Use this command to display existing scheduler configuration.

EXAMPLES

modify ltm classification auto-update settings auto-update-interval daily enabled yes

Enable hitless upgrade with update frequency set to daily.

modify ltm classification auto-update settings enabled no

Disable hitless upgrade.

list ltm classification auto-update

List the configured settings for auto-update settings configuration.

OPTIONS

enabled Specifies that the updates scheduler is enabled or disabled.

auto-update-interval

Specifies the auto-update frequency for classification signatures. This attribute will only apply in case auto update is enabled. The default value is weekly.

SEE ALSO

list, modify, load, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2018. All rights reserved.

BIG-IP 2018-03-16 ltm classification auto-update settings(1)

ltm classification auto-update status

NAME

status - Display the status for classification hitless upgrade.

MODULE

ltm classification classification auto-update status

SYNTAX

DISPLAY

list ltm classification auto-update status

options:

last-updated-time

message

progress-status

DESCRIPTION

Use this command to display current hitless upgrade status.

EXAMPLES

list ltm classification auto-update status

Display the status information for hitless upgrade.

OPTIONS

last-updated-time

Indicates the date and time of the last hitless upgrade attempt was made.

message

Indicates the success or error message after we attempt hitless upgrade.

progress-status

Indicates the progress status during hitless upgrade process.

SEE ALSO

list, ltm, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2018. All rights reserved.

BIG-IP 2018-03-16 ltm classification auto-update status(1)

Itm classification category

NAME

category - Configures a custom classification category.

MODULE

ltm classification

SYNTAX

Configure the category within the ltm classification module using the syntax shown in the following sections.

CREATE/MODIFY

create category [name]

options:

app-service [[string] | none]

description [string]

modify category [name]

options:

app-service [[string] | none]

description [string]

edit category [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list category

list category [[all] | [name]]

show running-config category

show running-config category [[all] | [name]]

options:

all-properties

non-default-properties

one-line
partition

DELETE
delete category [name]

Note: All referring applications/classification-filters (to this category) need to be deleted first; otherwise an error will be reported. Predefined categories cannot be deleted.

DESCRIPTION
You can use the category component to create, modify, delete, and display classification category.

EXAMPLES
create category my_cat

Creates a new category named my_cat.

```
modify category my_cat { description "My description." }
```

Modify a category named my_cat.

```
list category
```

Displays all created categories.

```
delete category my_cat
```

Deletes the category named my_cat.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description
User defined description.

SEE ALSO
create, modify, delete, list, show, tmsh, ltm classification

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2016-11-10 ltm classification category(1)

Itm classification ce

NAME
CE - Classification Engine configuration.

MODULE
ltm classification

SYNTAX
Configure the Classification Engine settings (AKA preset) within the ltm classification module using the syntax shown in the following sections.

CREATE/MODIFY
create ce [name]
modify ce [name]
options:
app-service [[string] | none]
allow-reclassification [on | off]
analyze-dns [on | off]
analyze-ssl-serverside [on | off]
flow-bundling [on | off]
cache-results [on | off]
policies [add | delete | default | replace-all-with | none]

DISPLAY
list ce [name]
show running-config ce [name]
options:
all-properties

non-default-properties
one-line

DELETE
delete ce [name]

DESCRIPTION
Use this command to create, modify, display, or delete an Classification Engine configuration.

EXAMPLES
create ce my_ce { allow-reclassification on flow-bundling on cache-results off }

Creates a new Classification Engine configuration named my_ce.

modify ce my_ce { cache-results on allow-reclassification off }

Turn Cache Results on and turn flow reclassification off on my_ce preset

list ce

Displays all Classification Engines configurations.

delete ce my_ce

Deletes the Classification Engine configuration named my_ce.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description
User defined description.

allow-reclassification
Specifies that connection ("flow") allowed to be reclassified, i.e. classification result could be changed.

analyze-dns
Allows classification engine to perform extended analysis of DNS traffic providing extended classification result.

analyze-ssl-serverside
Enables classification engine to process SSL Server Side Hello to inspect ALPN (mostly for HTTP2 / SPDY subclassification).

flow-bundling
Enables / disables flow correlation mechanism (e.g. FTP and FTP-data protocols).

cache-results
Enables / Disables classification result caching.

policies
Specifies a LTM policy that you have configured previously, if you want to configure custom signatures.

SEE ALSO
create, modify, delete, list, show, tmsh, ltm profile classification, ltm policy

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013, 2015-2016. All rights reserved.

BIG-IP 2016-10-05 ltm classification ce(1)

Itm classification signature-definition

NAME
signature-definition - Configure status for classification signature updates.

MODULE
ltm classification

SYNTAX
Configure the signature-definition component within the ltm classification module using the syntax in the following sections.

DISPLAY

list signature-definition

list signature-definition [[name] | [glob] | [regex]]

options:

all-properties

non-default-properties

one-line

recursive

last-attempt-automatic-mode [enabled | disabled]

last-attempt-datetime [date]

last-attempt-user [string]

last-update-automatic-mode [enabled | disabled]

last-update-datetime [date]

last-update-user [string]

message [string]

progress-status [none | success | failure | in-progress]

DESCRIPTION

You can use the signature-definition component to configure the status for classification signature updates.

EXAMPLES

```
list signature-definition
```

Displays classification signature update status configuration.

OPTIONS

last-attempt-automatic-mode

Indicates whether the last attempt to update the signature file was done manually or automatically by the system.

last-attempt-datetime

Indicates the date and time of the last attempt to update the signature file.

last-attempt-user

Indicates the user who is the last one to attempt to update the signature file.

last-update-automatic-mode

Indicates whether the last successful signature file update was done manually or automatically by the system. The value of the last-update-automatic-mode may be different from the value of the last-attempt-automatic-mode if the last update attempt fails.

last-update-datetime

Indicates the date and time of the last successful signature file update. The value of the last-update-datetime is different from the value of the last-attempt-datetime if the last update attempt fails.

last-update-user

Indicates the user who did the last successful signature file update. The value of the last-update-user may be different from the value of the last-attempt-user if the last update attempt fails.

message

Indicates the error message when it fails to attempt to update the signature file.

progress-status

Indicates the progress status when attempting to update the signature file. The options are none, success, failure, and in-progress.

SEE ALSO

list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-02-04 Itm classification signature-definition(1)

Itm classification signature-update-schedule

NAME

signature-update-schedule - Configure schedule for classification signature updates.

MODULE

itm classification

SYNTAX

Configure the signature-update-schedule component within the Itm classification module using the syntax in the

following sections.

Updated Command

There has been a new implementation for this command. Please consider using the following command instead, `ltm classification auto-update`

MODIFY

```
modify signature-update-schedule [name]
options:
  [auto-update-enabled | auto-update-disabled]
  auto-update-interval [daily | weekly | monthly]
```

```
edit signature-update-schedule [name]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list signature-update-schedule
list signature-update-schedule [ [name] | [glob] | [regex] ]
options:
  all-properties
  non-default-properties
  one-line
  recursive
```

DESCRIPTION

You can use the `signature-update-schedule` component to configure schedule for classification signature updates.

EXAMPLES

```
list signature-update-schedule
```

Displays classification signature update schedule configuration.

```
modify signature-update-schedule auto-update-enabled auto-update-interval daily
```

Updates the scheduler for classification signature updates to run once a day.

```
modify signature-update-schedule auto-update-disabled
```

Disables the scheduler and allows signatures to update via the browser-based Configuration utility only.

OPTIONS

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`auto-update-disabled`

Specifies that the updates scheduler is disabled. The user can update the classification signatures using the browser-based BIG-IP Configuration utility.

`auto-update-enabled`

Specifies that the updates scheduler is enabled.

`auto-update-interval`

Specifies the auto-update frequency for classification signatures. This attribute will only apply in case auto update is enabled. The default value is weekly.

SEE ALSO

`list`, `modify`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2018-02-05 ltm classification signature-update-schedule(1)

Itm classification signature-version

NAME

`signature-version` - Display classification signature version.

MODULE

ltm classification

SYNTAX

Display the signature-version component within the ltm classification module using the syntax in the following sections.

DISPLAY

list signature-version

options:

all-properties

non-default-properties

one-line

recursive

DESCRIPTION

You can use the signature-version component to display versions in classification signature updates.

EXAMPLES

list signature-version

Displays classification signature version configuration.

OPTIONS

cec-filename

Indicates the cec library filename in the last updated classification signature.

cec-version

Indicates the cec library version in the last updated classification signature.

classification-version

Indicates the classification version in the last updated classification signature.

conf-filename

Indicates the configuration filename in the last updated classification signature.

conf-version

Indicates the configuration version in the last updated classification signature.

im-version

Indicates the im version in the last updated classification signature.

qm-protocols-filename

Indicates the qosmos protocols filename in the last updated classification signature.

qm-protocols-version

Indicates the qosmos protocols version in the last updated classification signature.

update-time

Indicates the update time in the last updated classification signature.

SEE ALSO

list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2013-11-06 ltm classification signature-version(1)

ltm classification signatures

NAME

signatures - Manages classification engine signatures.

MODULE

ltm classification

SYNTAX

load signatures file [filename]

Loads classification signatures from a signature update file.

load signatures default

Resets classification engine and signatures to the factory defaults.

DESCRIPTION

You can use the signatures component to load the classification signatures from a file. Only admins can run this command.

You can obtain the latest signature update file (*.im) (if one is available) from <http://downloads.f5.com>.

For the filename, if no absolute path is specified, the default path /var/libdata/dpi/im/ is used.

Use load signatures default to discard and remove any installed upgrades and reset the classification engine and signatures to factory defaults. No user-created signatures will be deleted. By keeping the *.im file, you will be able to re-apply the update any time later using the load signatures file command.

EXAMPLES

```
load signatures file my_sig_file.im
```

Loads signatures from file "my_sig_file.im" under the folder: /var/libdata/dpi/im/.

```
load signatures file /var/tmp/new_sig_file.im
```

Loads signatures from file "new_sig_file.im" under the folder: /var/tmp/.

SEE ALSO

ltm classification signature-update-schedule, load, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2014-12-05 ltm classification signatures(1)

Itm classification stats application

NAME

application - Displays and resets classified application statistics.

MODULE

ltm classification stats

SYNTAX

Display statistics for the application component within the ltm classification stats module using the syntax in the following section.

DISPLAY

```
show application
```

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the application component to display classification application statistics. The statistics details are described below:

Name Specifies the number of the classified application.

Count

Specifies a number of classified flows or transactions (in transaction mode) to specific application.

LTM Policy

Specifies the number of classification decisions by LTM Policy (cpm).

Classification Engine

Specifies the number of classification decisions by classification engine (CEC).

Qosmos iXengine

Specifies the number of classification decisions by Qosmos iXEngine (ixe).

Cache

Specifies the number of classification decisions by Results Cache.

URI Parameter

Specifies the number of classification decisions by evaluating HTTP URI query string classification parameter.

HTTP Header

Specifies the number of classification decisions by using HTTP classification header.

iRule

Specifies the number of classification decisions by iRule.

Bytes in
Specifies the bytes in of the classified application.

Bytes out
Specifies the bytes out of the classified application.

Packets in
Specifies the packets in of the classified application.

Packets out
Specifies the packets out of the classified application.

You can reset the classification application statistics using reset-stats command.

EXAMPLES

show application

Displays the classified application statistics.

reset-stats application

Resets the classified application statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 Itm classification stats application(1)

Itm classification stats url-category

NAME

url-category - Displays and resets classified url-category statistics.

MODULE

Itm classification stats

SYNTAX

Display statistics for the url-category component within the Itm classification stats module using the syntax in the following section.

DISPLAY

show url-category

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the url-category component to display classification url-category statistics. The statistics details are described below:

Name Specifies the number of the classified url-category.

Count

Specifies a number of classified flows or transactions (in transaction mode) to specific url-category.

iRule

Specifies the number of url-categorization decision by iRule.

Customdb

Specifies the number of url-categorization decision by custom url database.

wrdb Specifies the number of url-categorization decision by webroot database.

Bytes in

Specifies the bytes in of the classified url-category.

Bytes out

Specifies the bytes out of the classified url-category.

Packets in
Specifies the packets in of the classified url-category.

Packets out
Specifies the packets out of the classified url-category.

You can reset the classification url-category statistics using reset-stats command.

EXAMPLES

show url-category

Displays the classified url-category statistics.

reset-stats url-category

Resets the classified url-category statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm classification stats url-category(1)

Itm classification stats urlcat-cloud

NAME

urlcat-cloud - Displays and resets URL category cloud lookup statistics.

MODULE

ltm classification stats

SYNTAX

Display statistics for the url category lookup within the ltm classification stats module using the syntax in the following section.

DISPLAY

show urlcat-cloud

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the urlcat-cloud component to display URL category cloud lookup statistics. The statistics details are described below:

Lookups

Specifies the number of URL look-up happened in cloud cache.

Cloud Query

Specifies the number of queries sent to the cloud. This happens if the entry is not present in cloud cache.

Cloud Response

Specifies the number of responses received from the cloud lookup. Each TMM gets a response. Hence the total is multiplied by TMM count.

Unknown

Specifies the number of unknown responses from the cloud. This means the webroot cloud does not know about the queried URL.

Entries

Specifies the number of entries present in the cloud cache. Each TMM has a copy of the cache. This number is cumulative and reset-stats does not reset this number.

Add Specifies the number of entries added to the cloud cache. If the cache is not cleared "adds" and "entries" must be the same.

Delete

Specifies the number of entries cleared from the cloud cache.

Queue Full

Specifies the number of instances that tmm could not send out a URL query.

You can reset the URL category cloud lookup statistics using reset-stats command. Reset-stats does not reset the number of entries in the cache.

EXAMPLES

show urlcat-cloud

Displays the URL category cloud lookup statistics.

reset-stats urlcat-cloud

Resets the URL category cloud lookup statistics, except the number of entries in the cache.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2016. All rights reserved.

BIG-IP 2016-03-14 Itm classification stats urlcat-cloud(1)

Itm classification update-signatures

NAME

update-signatures - Run automatic update for classification signatures.

MODULE

Itm classification

SYNTAX

Run the update-signatures component within the Itm classification module using the syntax in the following sections:

Updated Command

There has been a new implementation for this command. Please consider using the following command instead, run Itm classification updates

RUN

run update-signatures

DESCRIPTION

You can use the update-signatures component to update the classification signatures from F5 download server. Only admins can run this command.

SEE ALSO

Itm classification signature-update-schedule, run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2018-02-06 Itm classification update-signatures(1)

Itm classification updates

NAME

updates - Configures classification updates.

MODULE

Itm classification updates

SYNTAX

INSTALL

install ltm classification updates

options:

file [string]

LOAD

load ltm classification updates

options:

[string]

RUN

run ltm classification updates

DISPLAY

show ltm classification updates

DESCRIPTION

Use this command to install and see all available IM packages.

EXAMPLES

install ltm classification updates file file_name

Install new update from file "file_name".

load ltm classification updates file_name

Load new update from file "file_name".

run ltm classification updates

Downloads new classification update from staging or production download server.

show ltm classification updates

Show all available updates.

OPTIONS

file Specifies IM package file.

SEE ALSO

show, install, load, run, ltm, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2018. All rights reserved.

BIG-IP 2018-03-16 ltm classification updates(1)

Itm classification url-cat-policy

NAME

url-cat-policy - Configures an ltm classification url-cat policy. It's comprised of list of urldb feed lists.

This is deprecated since version 13.0.0. Instead of creating url-cat-policy and attaching a feedlist to it, enable the feedlist directly.

MODULE

ltm classification

SYNTAX

Configure the url-cat-policy component within the ltm classification module using the syntax in the following sections.

CREATE/MODIFY

create url-cat-policy [name]

modify url-cat-policy [name]

options:

app-service [name]

description [string]

feed-lists [add | delete] { [name] }

edit url-cat-policy

options:

all-properties

non-default-properties

DISPLAY
list url-cat-policy [[[name]]
show running-config url-cat-policy
show running-config url-cat-policy [[[name]]
options:
 all-properties
 non-default-properties
 one-line
 partition
 recursive

DISPLAY
delete url-cat-policy [name]

DESCRIPTION

You can use the url-cat-policy component to configure a shareable and reusable url categorization database feed coming from local files or download feeds. The url-cat-policy can then be enforced on the configuration object of the type: ltm virtual.

EXAMPLES

```
create ltm classification url-cat-policy POL1
feed-lists add { FL1 }
description none }
```

Creates a url-cat-policy POL1 with urldb feeds from FL1 feed lists.

```
list url-cat-policy
```

Displays the current list of url-categorization policies contents.

OPTIONS

app-service
Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description
User defined description.

partition
Displays the administrative partition within which the component resides.

feed-lists
Adds, deletes, or replaces a feed list. Specifies a list of feed lists (see ltm classification urldb-feed-list) against which the packet will be compared.

SEE ALSO

create, edit, list, modify, ltm classification urldb-feed-list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2015. All rights reserved.

BIG-IP 2016-10-28 ltm classification url-cat-policy(1)

ltm classification url-category

NAME
url-category - Configures a custom URL category.

MODULE
ltm classification

SYNTAX
Configure the url-category within the ltm classification module using the syntax shown in the following sections.

CREATE/MODIFY
create url-category [name]
options:
 app-service [[string] | none]
 description [string]
 url-category-id [integer]
 irule-event [enabled | disabled]

modify url-category [name]

options:
app-service [[string] | none]
description [string]
irule-event [enabled | disabled]

edit url-category [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list url-category
list url-category [[all] | [name]]
show running-config url-category
show running-config url-category [[all] | [name]]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete url-category [name]

Note: All referring url-categorization-filters (to this url-category) need to be deleted first; otherwise an error will be reported. Predefined url-categories cannot be deleted.

DESCRIPTION

You can use the url-category component to create, modify, delete, and display classification url-category.

EXAMPLES

```
create url-category my_urlcat { url-category-id 28672 irule-event disabled }
```

Creates a new url-category named my_urlcat.

```
modify url-category my_urlcat { irule-event enabled description "My description." }
```

Modify a url-category named my_urlcat.

```
list url-category
```

Displays all created categories.

```
delete url-category my_urlcat
```

Deletes the url-category named my_urlcat.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

url-category-id

Identifies the url-category. This is set during creation and cannot be changed. Identifiers must be unique across predefined and user-defined categories. Predefined url-category-ids must be in numeric range [24576, 28671), and user defined url-category-ids must be in numeric range [28672-32768).

irule-event

Indicates if the irule is enabled or disabled in result of the classification engine.

SEE ALSO

create, modify, delete, list, show, tmsh, ltm classification, pem policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013. All rights reserved.

BIG-IP 2013-11-12 ltm classification url-category(1)

ltm classification urldb-feed-list

NAME

urldb-feed-list - Configures a feed-list to be used for URLDB file loads. A urldb-feed-list is a list of URL feeds (including local file paths) from where URLDB files are downloaded. These files contain URL categorization information.

MODULE

ltm classification

SYNTAX

Configure the urldb-feed-list component within the ltm classification module using the syntax in the following sections.

CREATE/MODIFY

```
create urldb-feed-list [name]
modify urldb-feed-list [[name] | all]
options:
  default-url-category [name]
  url [string]
  poll-interval [integer]
  user [string]
  password [string]
  app-service [name]
  description [string]
```

```
load urldb-feed-list [[name] | all]
```

DISPLAY

```
list urldb-feed-list [[name] | all | [property]]
show running-config urldb-feed-list [[name] | all | [property]]
options:
  all-properties
  non-default-properties
  one-line
  partition
  recursive
```

DELETE

```
delete urldb-feed-list [[name] | all]
```

DESCRIPTION

You can use the urldb-feed-list component to define reusable lists of feeds. You can use a feed list in an ltm classification url-cat-policy.

EXAMPLES

```
create urldb-feed-list FL1 { url file:///shared/images/custom_urldb_1.txt }
```

Creates a new feed list, "FL1" with URL category information in the file specified by url.

OPTIONS

create
Creates a new feed list.

delete
Deletes the feed list that you specify next, in curly braces ({}).

file
DEPRECATED since version 11.7.0. Specifies the file object containing the URLDB information.

url Specifies the url to fetch the file containing the URLDB information.

default-url-category
The URL category to be used for all the URLs specified in the URLDB file.

poll-interval
Specifies the time interval in hours at which the url needs to be polled.

app-service
Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description
User defined description for this feed list.

partition
Displays the administrative partition within which the component resides.

SEE ALSO

edit, list, modify, ltm classification, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2014, 2016. All rights reserved.

Itm classification urldb-file

NAME

urldb-file - Manages a custom URLDB file

MODULE

Itm classification

SYNTAX

List the urldb-file component within the Itm classification module using the syntax in the following sections.

CREATE

create urldb-file [name]

options:

source-path [string]

app-service [name]

DEPRECATED: create command is deprecated from version 12.0.0. Though this command is visible, this is not meant to be used by the users. The daemons use it internally.

DISPLAY

list urldb-file [[name] | all | [property]]

show running-config urldb-file [[name] | all | [property]]

options:

all-properties

one-line

partition

DELETE

delete urldb-file [[name] | all]

DEPRECATED: delete command is deprecated from version 12.0.0. Though this command is visible, this is not meant to be used by the users. The daemons use it internally.

DESCRIPTION

You can use the urldb-file component to create, the custom URLDB file in versions before 11.7.0. The urldb-file is created internally using the URI specified in Itm classification urldb-feed-list in later versions.

EXAMPLES

```
create urldb-file FILE1 { source-path file:/shared/images/custom_urldb_1.txt }
```

Creates a new urldb file object, "FILE1" from the source file /shared/images/custom_urldb_1.txt

Create works in versions before 11.7.0. Has been deprecated in later versions

```
list urldb-file [fileobj-name]
```

Lists the attributes of urldb file object, "FILE1" from the source file /shared/images/custom_urldb_1.txt

OPTIONS

create

Creates a new file object for custom urlcat db.

delete

Deletes the file object that you specify next, in curly braces ({}).

source-path

Specifies the location from where the URLDB file object sources the file.

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description

User defined description for this feed list.

partition

Displays the administrative partition within which the component resides.

SEE ALSO

edit, list, Itm classification, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

Itm clientssl-proxy cached-certs

NAME

cached-certs - Displays and deletes SSL Forward Proxy cached certificates and OCSP responses on the BIG-IP(r) system.

MODULE

ltm clientssl-proxy

SYNTAX

Use the cached-certs component within the ltm.clientssl-proxy module to manage connections using the following syntax.

DISPLAY

show cached-certs

options:

virtual [name]

clientssl-profile [name]

DELETE

delete cached-certs

options:

virtual [name]

clientssl-profile [name]

DESCRIPTION

You can use the cached-certs component to display or delete SSL Forward Proxy cached certificates based on a specified clientssl profile.

OPTIONS

virtual

Specifies the name of the virtual server that you want to display or delete cached certificates from.

clientssl-profile

Specifies the name of the clientssl profile that belongs to the virtual selected.

SEE ALSO

delete, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2014. All rights reserved.

Itm clientssl ocspp-stapling-responses

NAME

ocsp-stapling-responses - Deletes the cached OCSP responses on the BIG-IP(r) system.

MODULE

ltm clientssl

SYNTAX

Use the ocsp-stapling-responses component within the ltm.clientssl module to manage connections using the following syntax.

DELETE

delete ocsp-stapling-responses

options:

virtual [name]

clientssl-profile [name]

DESCRIPTION

You can use the ocsp-stapling-responses component to delete the cached OCSP responses based on a specified clientssl profile.

OPTIONS

virtual

Specifies the name of the virtual server that you want to display or delete cached certificates from.

clientssl-profile

Specifies the name of the clientssl profile that belongs to the virtual selected.

SEE ALSO

delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-05-30 ltm clientssl ocsp-stapling-responses(1)

ltm data-group external

NAME

external - Configures an external class.

MODULE

ltm data-group

SYNTAX

Configure the external data-group within the ltm data-group module using the syntax shown in the following sections.

CREATE/MODIFY

create external [name]

modify external [name]

options:

app-service [[string] | none]

description [string]

external-file-name [[file name] | none]

separator [string]

source-path [URL]

type [integer | ip | string]

edit external [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list external

list external [[[name] | [glob] | [regex]] ...]

show [name] external-records

show running-config external

show running-config external [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete external [name]

DESCRIPTION

Data groups are lists of data that you define and use with iRules(r) operators. External data group records are stored in external files that you manage through the sys file data-group component. Note that external data groups can be very large, which is one reason why the groups are saved to external files. For example, a phone company may store a list of thousands of phone numbers in an external data group.

You should consider using an internal data group when the number of records is expected to be small.

An external data group acquires its type from the associated data-group file, which can be a list of IP addresses, strings, or integers.

External data groups are lists that specify:

• A data-group file where records are stored

• A description of the class

There are two ways to configure the external data-group object:

• Create external data-group object, and then specify the source-path and type of the external-file. In one step the external-file will be created within the sys file data-group module and external data-group within the ltm data-group module.

• Create an external-file within the sys file data-group module, and then create external data-group within the ltm data-group module. See help sys file data-group for information on creating the data-group file.

EXAMPLES

`create external ext-dg1 external-file-name string.dat description "created for rule xyz"`
Creates an external data group named ext-dg1, with the given description. The records for the data group are loaded from the data-group file string.dat previously created in the sys file data-group component.

`create external ext-dg1 description "created for rule xyz" source-path http://file-server/data-groups/ip.class type ip`
Downloads the data-group file from the given URL into file-store and creates a data-group file named ext-dg1 within the sys file data-group module. Creates an external data group named ext-dg1, with the given description. The records for the data group are loaded from the data-group file ext-dg1.

`create external ext-dg2 source-path file:/shared/save/Test.dat type string`
Creates a data-group file named ext-dg2 within the sys file data-group module. Creates an external data group named ext-dg2. The records for the data group are loaded from the data-group file ext-dg2.

`modify external ext-dg2 description "created for rule abc" source-path file:/shared/save/Test2.dat`
Downloads the file from the given URL into file-store and updates the source-path of data-group file referenced by external data group ext-dg2. Modifies the description of external data group ext-dg2.

`delete external ext-dg1`
Deletes the external data group named ext-dg1. Note: the data-group file referenced by ext-dg1 is not deleted at this time. If needed, it should be deleted under sys file data-group component.

`show external ext-dg1 external-records` Shows a sorted list of external data group records. Note: We recommend against displaying data-groups of over two million records due to control-plane resource constraints.

OPTIONS

`app-service`
Specifies the name of the application service to which the data group belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the data group. Only the application service can modify or delete the data group.

`description`
User defined description.

`external-file-name`
Specifies the data-group file where the records are stored.

Note: Only source-path or external-file-name may be specified for external data-group configuration item.

`regex`
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`separator`
Specifies a separator to use when defining the data group. The default value is :=.

`source-path [URL]`
This optional attribute takes a URL.

Note: Only source-path or external-file-name may be specified for external data-group configuration item, for example:

`source-path http://file-server/data-groups/AUL_1.cls`

`source-path https://file-server/data-groups/CNN.x`

`source-path ftp://username:password@server/data-groups/latest.class`

`source-path file:/shared/save/Test.dat`

`type` Specifies the kind of data in the group. This option is acquired from the data group file. If the external data group is created with external-file that was previously created within the sys file data-group module, then type option cannot be modified. If the external data group is created with source-path option, then type should be specified. The value for type could be integer or ip or string.

SEE ALSO

create, delete, edit, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

ltm data-group internal

NAME

internal - Configures an internal class.

MODULE

ltm data-group

SYNTAX

Configure the internal data-group within the ltm data-group module using the syntax shown in the following sections.

CREATE/MODIFY

create internal [name]

options:

app-service [[string] | none]

description [string]

records [add | delete | modify | replace-all-with] {

[record key] {

data [value]

}

records none

type [integer | ip | [string]]

modify internal [name]

options:

app-service [[string] | none]

description [string]

records [add | delete | modify | replace-all-with] {

[record key] {

data [value]

}

records none

edit internal [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list internal

list internal [[name] | [glob] | [regex]] ...]

show running-config internal

show running-config internal [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DESCRIPTION

Data groups are lists of data that you define and use with iRules(r) operators. Consider using an external data group if the number of records is expected to be large.

The BIG-IP(r) system includes a number of predefined lists that you can use. They are:

• aol

• default_accept_language

• images

• private_net

The above lists are located in the file /config/profile_base.conf. When you run the load command, the system loads these lists; however, unless you have modified the lists, the system does not save the lists to the bigip.conf file.

The internal data groups are stored in the bigip.conf file.

Internal data groups can be one of three types:

• A list of IP addresses

• A list of strings

Â· A list of integers

Strings must be surrounded by quotation marks. Numbers can be either positive or negative. These groups define the type of data in the class, which can be IP addresses, strings, or integers>

EXAMPLES

```
create internal MyDG records add { 10.0.0.0 } type ip
```

Creates an internal data group named MyDG that contains a single IP address.

```
create internal DG2 records add { 192.1.1.255 192.2.1.255 192.3.1.255 } type ip
```

Creates an internal data group named DG2 that contains a list of three network addresses: 192.1.1.0/24, 192.2.1.1/24, and 192.3.1.1/24.

```
create internal MyDG records add { my_key { data my_value } } type string
```

Creates an internal data group named MyDG that contains a single name/value pair.

OPTIONS

app-service

Specifies the name of the application service to which the data group belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the data group. Only the application service can modify or delete the data group.

description

User defined description.

records

Configures the data in the group.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

type Specifies the kind of data in the group. The default value is ip. This option is required by the command create.

SEE ALSO

create, delete, edit, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-12-21 ltm data-group internal(1)

ltm default-node-monitor

NAME

default-node-monitor - Configures the default node monitor for the Local Traffic Manager.

MODULE

ltm

SYNTAX

Configure the default-node-monitor component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

```
modify default-node-monitor
options:
rule [rule syntax]
```

```
edit default-node-monitor
```

```
options:
all-properties
```

DISPLAY

```
list default-node-monitor
show running-config default-node-monitor
options:
one-line
all-properties
```

DESCRIPTION

You can use the default-node-monitor component to modify the default monitor that the system applies to any node address to which a monitor is not explicitly assigned.

EXAMPLES

```
modify default-node-monitor rule icmp
```

Modifies the global default node monitor to use the rule ICMP.

```
list default-node-monitor
```

Displays the properties of the global default node monitor.

OPTIONS

rule Specifies the rule that the system applies to any node that has not been assigned a monitor rule. The default value is none.

You can specify:

Â· A single monitor, for example, modify default-node-monitor rule icmp.

Â· Multiple monitors, for example, modify default-node-monitor rule icmp and tcp_echo.

Â· A minimum number of monitors, for example, modify default-node-monitor rule min 1 of { icmp and tcp_echo }.

SEE ALSO

list, ltm node, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2013-08-30 ltm default-node-monitor(1)

ltm dns analytics global-settings

NAME

global-settings - Configures the global settings of all DNS listeners on the BIG-IP(r) system.

MODULE

ltm dns analytics

SYNTAX

Configure the global-settings DNS listeners within the ltm dns analytics module using the syntax in the following sections.

MODIFY

```
modify global-settings
```

options:

```
collect-client-ip [enabled | disabled]
```

```
collect-query-name [enabled | disabled]
```

DISPLAY

```
list global-settings
```

```
list global-settings
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DESCRIPTION

You can use the global-settings component to configure and view information about the global settings of all DNS listeners.

EXAMPLES

```
list global-settings all-properties
```

Displays the global settings for the DNS listeners on the BIG-IP system.

OPTIONS

collect-client-ip

When enabled, the client IP addresses of DNS queries will be collected and stored in analytics database. The default value is enabled.

collect-query-name

When enabled, the domain names of DNS queries will be collected and stored in analytics database. The default value is enabled.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-06-28 ltm dns analytics global-settings(1)

ltm dns cache global-settings

NAME

global-settings - Configures the global settings of all DNS caches on the BIG-IP(r) system.

MODULE

ltm dns cache

SYNTAX

Configure the global-settings DNS cache component within the ltm dns cache module using the syntax in the following sections.

CREATE/MODIFY

modify global-settings [name]

options:

cache-maximum-ttl [integer]

cache-minimum-ttl [integer]

resolver-edns-buffer-size [integer]

edit global-settings [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list global-settings

list global-settings [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the global-settings component to configure and view information about the global settings of all DNS caches.

EXAMPLES

list global-settings all-properties

Displays the global settings for the DNS caches on the BIG-IP system.

OPTIONS

cache-maximum-ttl

Specifies the number of seconds after which you want the BIG-IP system to re-query for resource records. This setting allows the BIG-IP system to re-query for resource records sooner than the owner of the records intended.

cache-minimum-ttl

Specifies the minimum number of seconds you want the BIG-IP system to cache DNS resource records. This setting allows the BIG-IP system to cache resource records longer than the owner of the records intended.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

resolver-edns-buffer-size [integer]

Specifies the number of bytes you want the BIG-IP system to advertise as the EDNS buffer size in UDP queries.

SEE ALSO

edit, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-03-21 Itm dns cache global-settings(1)

Itm dns cache records all

NAME

all - Simultaneously manages all subcache types in DNS cache resolvers on the BIG-IP(r) system.

MODULE

itm dns cache records

SYNTAX

Configure all cache components within the itm dns cache records module using the syntax in the following sections.

DELETE

delete all cache [cache name]

DESCRIPTION

This all cache option clears all subcaches at once and resizes them to their initial sizes.

EXAMPLES

delete itm dns cache records all cache my_dns_cache_resolver

Deletes all entries of all subcaches in my_dns_cache_resolver.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc. F5 Networks and BIG-IP (c) Copyright 2009-2019. All rights reserved.

BIG-IP 2019-04-17 Itm dns cache records all(1)

Itm dns cache records key

NAME

key - Manages the DNSKEY records in the DNS caches on the BIG-IP(r) system.

MODULE

itm dns cache records

SYNTAX

Configure the key component within the itm dns cache records module using the following syntax.

DISPLAY

show key cache [cache name]

options:

count-only

owner [domain name]

slot [integer]

tmm [integer]

DELETE

delete key cache [cache name]

options:

owner [domain name]

EXAMPLES

show key cache resolver_cache

Displays the DNSKEY records in the cache named resolver_cache.

delete key cache v_resolver_cache

Deletes the DNSKEY records from the cache named v_resolver_cache.

DESCRIPTION

You can use the following options with the key component.

OPTIONS

cache name

Specifies a DNS cache name from which to display or delete DNSKEY records. This is a required field.

count-only

For a show command, return only a count of the number of matched records.

owner

Specifies a domain name on which to filter the DNSKEY records in the specified DNS cache for a query or deletion.

slot Specifies a slot number on a chassis that contains the specified DNS cache. This is a 1 based index.

tmm Deprecated in v15.0.0 Use tmm-process instead. All records across all TMM processes and slots will be returned.

tmm-process

Specifies which individual TMM process's key records you want to view. This is a 0 based index. On a chassis, this filter also requires a slot to be specified. If not specified, all records will be returned unless the slot filter is in use.

SEE ALSO

delete, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2019-04-17 ltm dns cache records key(1)

ltm dns cache records msg

NAME

msg - Manages message records in the DNS caches on the BIG-IP(r) system.

MODULE

ltm dns cache records

SYNTAX

Configure the msg component within the ltm dns cache records module using the following syntax.

DISPLAY

show msg cache [cache name]

options:

count-only

qname [domain name]

rcode [integer]

slot [integer]

tmm [integer]

DELETE

delete msg cache [cache name]

options:

qname [domain name]

rcode [integer]

DESCRIPTION

The msg component contains full DNS messages. You can display and delete these messages.

EXAMPLES

show msg cache resolver_cache

Displays the message records in the DNS cache named resolver_cache.

delete msg cache v_resolver_cache

Deletes the message records from the DNS cache named v_resolver_cache.

OPTIONS

cache name

Specifies a DNS cache name. This is a required field.

count-only

For a show command, return only a count of the number of matched records.

qname

Specifies a domain name on which to filter the DNS messages in the specified DNS cache for a query or deletion.

rcode

Specifies the DNS return code on which to filter DNS messages in the specified DNS cache for a query or deletion.

slot Specifies a slot number on a chassis that contains the specified DNS cache. This is a 1 based index.

tmm Deprecated in v15.0.0 Use tmm-process instead. All records across all TMM processes and slots will be returned.

tmm-process

Specifies which individual TMM process's msg records you want to view. This is a 0 based index. On a chassis, this filter also requires a slot to be specified. If not specified, all records will be returned unless the slot filter is in use.

SEE ALSO

delete, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2019-04-17 ltm dns cache records msg(1)

ltm dns cache records nameserver

NAME

nameserver - Manages the nameserver records in the DNS cache resolvers on the BIG-IP(r) system.

MODULE

ltm dns cache records

SYNTAX

Configure the nameserver component within the ltm dns cache records

DISPLAY

show cache [cache name]

options:

address [ip address]

count-only

has-edns [yes | no]

has-lame [yes | no]

rtt-range [min:max]

slot [integer]

tmm [integer]

ttl-range [min:max]

zone-name [name]

DELETE

delete cache [cache name]

options:

address [ip address]

has-edns [yes | no]

has-lame [yes | no]

rtt-range [min:max]

ttl-range [min:max]

zone-name [name]

DESCRIPTION

You can use the nameserver component to display or delete nameserver records from a DNS cache. The maximum number of records returned is 1000; therefore, broad searches may not show all records in the cache.

EXAMPLES

show cache my_cache zone-name com ttl-range 50:500

Displays the nameserver records, in the DNS cache named my_cache, with the zone name com, where the TTLs of the records are between 50 and 500.

OPTIONS

address

Specifies the nameserver records, in the specified DNS cache, to select based on the IP address of the nameserver.

cache name

Specifies a DNS cache name. This is a required field.

count-only

For a show command, return only a count of the matched records.

has-edns

Specifies the nameserver records to select from the specified DNS cache, based on whether the nameserver is EDNS lame. An EDNS lame nameserver does not reply to EDNS queries.

has-lame

Specifies the nameserver records to select from the specified DNS cache, based on whether the nameserver is lame for one or more items.

rtt-range

Specifies the nameserver records to select from the specified DNS cache, based on RTTs within the specified range (inclusive). A missing value (:500 or 50:) defaults to the minimum or maximum, respectively.

slot Specifies a slot number on a chassis that contains the specified DNS cache. This is a 1 based index.

tmm Deprecated in v15.0.0 Use tmm-process instead. All records across all TMM processes and slots will be returned.

tmm-process

Specifies which individual TMM process's nameserver records you want to view. This is a 0 based index. On a chassis, this filter also requires a slot to be specified. If not specified, all records will be returned unless the slot filter is in use.

ttl-range

Specifies the nameserver records to select from the specified DNS cache, based on TTLs within the specified range (inclusive). A missing value (:500 or 50:) defaults to the minimum or maximum, respectively.

zone-name

Specifies the nameserver records to select from the specified DNS cache, based on the specified zone name.

SEE ALSO

delete, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2019-04-17 ltm dns cache records nameserver(1)

Itn dns cache records rrset

NAME

rrset - Manages the RRset records in the DNS cache resolvers on the BIG-IP(r) system.

MODULE

ltm dns cache records

SYNTAX

Configure the rrset component within the ltm dns cache records module using the syntax in the following sections.

DISPLAY

show cache [cache name]

options:

class [IN | CH | HS| ANY]

count-only

owner [DNS name]

slot [integer]

tmm [integer]

ttl-range [integer:integer]

type [A | AAAA | CNAME | NS | PTR | RRSIG | DNSKEY | SOA | TXT | ANY | ...]

DELETE

delete cache [cache name]

options:

class [IN | CH | HS | ANY]

owner [DNS name]

ttl-range [integer:integer]

type [A | AAAA | CNAME | NS | PTR | RRSIG | DNSKEY | SOA | TXT | ANY | ...]

DESCRIPTION

You can use the rrset component to display or delete records in the specified DNS cache. The maximum number of records returned is 1000. Broad searches might not show all records in the cache.

EXAMPLES

```
show cache resCache2 class IN type A ttl-range 20:5000 owner .com
```

Displays RRset records of type A, class IN, with TTLs between 20 and 5000, and an owner of .com.

OPTIONS

cache name

Specifies a DNS cache name. This is a required field.

class

Specifies the class of RRset records to select from the specified DNS cache.

count-only

For a show command, return only a count of the matched records.

owner

Specifies the node on which to filter the RRset records in the specified DNS cache for a query or deletion.

slot Specifies a slot number on a chassis that contains the specified DNS cache. This is a 1 based index.

tmm Deprecated in v15.0.0 Use tmm-process instead. All records across all TMM processes and slots will be returned.

tmm-process

Specifies which individual TMM process's rrset records you want to view. This is a 0 based index. On a chassis, this filter also requires a slot to be specified. If not specified, all records will be returned unless the slot filter is in use.

ttl-range

Specifies the RRset records to select from the specified DNS cache, based on TTLs within the specified range (inclusive). A missing value (:500 or 50:) defaults to the minimum or maximum, respectively.

type Specifies the RRset records to select from the specified DNS cache, based on the specified type. Most record types are supported.

SEE ALSO

show, delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2019-04-17 Itm dns cache records rrset(1)

Itm dns cache resolver

NAME

resolver - Configures a DNS cache with a resolver on the BIG-IP(r) system.

MODULE

itm dns cache

SYNTAX

Configure the resolver DNS cache component within the itm dns cache module using the syntax in the following sections.

CREATE/MODIFY

```
create resolver [name]
```

```
modify resolver [name]
```

options:

allowed-query-time [integer]

answer-default-zones [yes | no]

app-service [[string] | none]

```

description [[string] | none]
forward-zones [add | delete | modify | replace-all-with] {
  [ [zone-name] ] {
options:
nameservers [add | delete | replace-all-with] {
  [ [IPv4address:port] | [IPv6address.port] ]
}
nameservers none
}
forward-zones none
local-zones [ [none] ]
[ { { name [dname] type [type] records [none | add { [RR string] ... } ] ... } } ]
max-concurrent-queries [integer]
max-concurrent-tcp [integer]
max-concurrent-udp [integer]
msg-cache-size [integer]
nameserver-cache-count [integer]
randomize-query-name-case [yes | no]
response-policy-zones [add | delete | modify] {
  [zone-name] {
action [nxdomain | walled-garden]
walled-garden [local-zone]
}
}
response-policy-zones none
root-hints {
  { [IP address] ... }
}
route-domain [name]
rrset-cache-size [integer]
rrset-rotate [none | query-id]
unwanted-query-reply-threshold [integer]
use-ipv4 [yes | no]
use-ipv6 [yes | no]
use-tcp [yes | no]
use-udp [yes | no]

```

DISPLAY

```

list resolver
list resolver [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
show resolver
show resolver [name]

```

DELETE

```

delete resolver [name]

```

DESCRIPTION

You can use the resolver component to configure and view information about a recursive-resolving DNS cache. A resolver cache performs recursive resolution to fill its cache.

Important: When sizing caches, consider the total amount of memory available and how you wish to allocate memory for DNS caching. Note that cache sizing values are per-TMM process; therefore, a platform with eight TMMs consumes the amount of memory set for the RRset cache times eight.

EXAMPLES

```

list resolver myCache

```

Displays the properties of the recursive-resolving DNS cache myCache.

```

modify resolver myCache local-zones { { name lz.example.net records add { "lz.example.net 60 IN A 127.0.0.1"
"www.lz.example.net 300 IN A 127.0.0.2" } } }

```

Modifies DNS cache myCache by adding a local-zone lz.example.net with 2 resource records.

OPTIONS

allowed-query-time

Specifies the time allowed for a query to stay in the queue before it is replaced by a new query when the number of concurrent distinct queries exceeds the limit. The default value is 200 milliseconds.

answer-default-zones

Specifies whether the resolver cache answers queries for default zones: localhost, reverse 127.0.0.1 and ::1, and AS112 zones. The default value is no.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

forward-zones

Adds, deletes, modifies, or replaces a set of forward zones on a DNS Cache, by specifying zone name(s). A given zone name should only use the symbols allowed for a fully qualified domain name (FQDN), namely ASCII letters a through z, digits 0 through 9, hyphen -, and period .. For example site.example.com would be a valid zone name.

A DNS Cache configured with a forward zone will forward any queries that result in a cache-miss (the answer was not available in the cache) and match a configured zone name, to the nameserver specified on the zone. If no nameservers are specified on the zone, an automatic SERVFAIL is returned. When a forward zone's nameserver returns a valid response to the DNS Cache, that response is cached and then returned to the requester.

nameservers

Adds, deletes, modifies, or replaces a set of nameservers in a forward zone on a DNS Cache. A nameserver is represented by an IP address and port in the format [IPv4:port] or [IPv6.port], for example 10.10.10.10:53 or 2001::1:ff.53, respectively.

If more than one nameserver is listed for a given forward zone, a matching query will be sent to the nameserver that is currently deemed the most responsive (based on RTTs). If no response is received within a certain window of time, the DNS Cache will resend the query to another nameserver with an increased wait window until a response is received.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

local-zones

Zones and associated resource records for which the cache will provide Authoritative responses. Default is empty. This is intended for small, simple authoritative data configurations.

The local-zone name must be fully qualified and should be the apex of the zone. The local-zone type may be one of the following: deny, refuse, static, transparent, type-transparent, or redirect. Zero or more resource records must be fully specified: name, ttl, class, type, and record data, separated by spaces, and within double quotes. For example, "www.example.net. 300 IN A 1.2.3.4".

For all local-zones types, if the DNS query matches, it is answered Authoritatively. How a non-matching query is handled depends on the local-zone type.

deny drops the query.

refuse sends a REFUSED response.

static sends either a NoData or NXDOMAIN response (includes SOA if present in local-zone).

transparent performs regular cache operation (i.e. transparent pass-through or iterative resolution) except for those query names which would result in NoData. This is the default local-zone type.

type-transparent Same as transparent but does not return NoData.

redirect returns responses with zone suffix record(s) for queries beneath that suffix. For example, a local-zone for example.com and a single A record for that name; queries for www.example.com or abc.www.example.com would return the single A record (both have the same suffix).

max-concurrent-queries

Specifies the maximum number of concurrent distinct queries used by the resolver. A query is identified by query name, type and class. If the number of distinct queries exceeds this limit, the resolver replaces the earliest query in the queue with the new query if it has been in the queue longer than the allowed time. The default value is 1024.

max-concurrent-tcp

Specifies the maximum number of concurrent TCP flows used by the resolver. The default value is 20.

max-concurrent-udp

Specifies the maximum number of concurrent UDP flows used by the resolver. The default value is 8192.

msg-cache-size

Specifies the maximum size in bytes of the DNS message cache. The default value is 1048576.

The BIG-IP system caches the messages in a DNS response in the message cache. After the maximum size of the cache is reached, when new or refreshed content is added to the cache, the expired and older content is removed from the cache. A higher maximum size allows more DNS responses to be cached and increases the cache hit percentage. A lower maximum size forces earlier eviction of cached content, but can lower the cache hit percentage.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

nameserver-cache-count

Specifies the maximum number of DNS nameservers for which the BIG-IP system caches connection and capability data. The default value is 16536 entries.

randomize-query-name-case

When enabled, the resolver randomizes the case of query names. The default value is yes.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

response-policy-zones

Adds, deletes or modifies the response policy zone to be used by this DNS Cache. Only a DNS Express zone configured as a response policy zone can be added.

The query name of a recursive DNS request without DNSSEC enabled is queried against the data in the response policy zone. If a match is found, the configured response policy action is taken.

action

The action to take upon a match. `nxdomain` results in an `NXDOMAIN` response given to the client. `walled-garden` results in a response with a `CNAME` to the `walled-garden` zone and an `A` or `AAAA` response matching the DNS query type. The default action is `nxdomain`.

walled-garden

A local zone configured in this cache that contains an `A` and/or `AAAA` record. This is typically used to redirect a user that requests resolution of a name contained in the RPZ database to a local server. This local server can display a message to the user and/or record the connection. Only `A/AAAA/ANY` requests are redirected, a request for any other type is answered with a `NoData` response. If a request is received for type `A` or `AAAA` but there are no records of that type configured, a `NoData` response is returned instead.

root-hints

Specifies the IP addresses of DNS servers that the BIG-IP system considers authoritative for the DNS root nameservers.

Important:By default, the BIG-IP system uses the DNS root nameservers published by InterNIC.

Caution:When you add DNS root nameservers, the BIG-IP system no longer uses the default nameservers published by InterNIC, but instead uses the nameservers you add as authoritative for the DNS root nameservers.

route-domain

Specifies the route domain the resolver uses for outbound traffic. The default value is the default route domain.

rrset-cache-size

Specifies the maximum size in bytes of the resource records set cache. The default value is 10485760.

The BIG-IP system caches the supporting records in a DNS response in the resource record cache. After the maximum size of the cache is reached, when new or refreshed content is added to the cache, the expired and older content is removed from the cache. A higher maximum size allows more DNS responses to be cached and increases the cache hit percentage. A lower maximum size forces earlier eviction of cached content, but can lower the cache hit percentage.

rrset-rotate

Specifies the resource record rotation method used within cached responses. The default value is `none`.

`none` Resource record order is not modified.

`query-id` Resource record order is a function of the client's query id.

unwanted-query-reply-threshold

The system always rejects unsolicited replies. The default value of 0 (`off`) indicates the system does not generate SNMP traps or log messages when rejecting unsolicited replies.

Change the default value to monitor for unsolicited replies. This alerts you to a potential security attack, such as cache poisoning or DOS. For example, if you specify a value of 1,000,000, each time the system receives 1,000,000 unsolicited replies, it generates an SNMP trap and log message. The default value is 0 (`off`).

use-ipv4

When enabled, the resolver sends DNS queries to IPv4 addresses. The default value is `yes`.

use-ipv6

When enabled, the resolver sends DNS queries to IPv6 addresses. The default value is `yes`.

use-tcp

When enabled, the resolver can send queries over the TCP protocol. The default value is `yes`.

use-udp

When enabled, the resolver can send queries over the UDP protocol. The default value is `yes`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm dns cache transparent`, `ltm dns cache validating-resolver`, `show`, `modify`, `regex`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2016. All rights reserved.

BIG-IP 2016-03-14 ltm dns cache resolver(1)

ltm dns cache transparent

NAME

transparent - Configures a DNS cache without a resolver on the BIG-IP(r) system.

MODULE

ltm dns cache

SYNTAX

Configure the transparent DNS cache component within the ltm dns cache module using the syntax in the following sections.

CREATE/MODIFY

```
create transparent [name]
modify transparent [name]
options:
  answer-default-zones [yes | no]
  app-service [[string] | none]
  description [[string] | none]
  local-zones [ [none] |
  [ { { name [dname] type [type] records [none | add { [RR string] ... } ] } ... } ] ]
  msg-cache-size [integer]
  response-policy-zones [add | delete | modify] {
    [zone-name] {
  action [nxdomain | walled-garden]
  walled-garden [local-zone]
    }
  }
  response-policy-zones none
  rrset-cache-size [integer]
  rrset-rotate [none | query-id]
```

DISPLAY

```
list transparent
list transparent [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
show transparent
show transparent [name]
```

DELETE

```
delete transparent [name]
```

DESCRIPTION

You can use the transparent component to configure and view information about a transparent DNS cache. A transparent cache does not perform recursive resolution, but instead relies on another DNS resource for this functionality.

Important: When sizing caches, consider the total amount of memory available and how you wish to allocate memory for DNS caching. Note that cache sizing values are per-TMM process; therefore, a platform with eight TMMs consumes the amount of memory set for the RRset cache times eight.

EXAMPLES

```
list transparent myCache
```

Displays the properties of the transparent DNS cache myCache.

```
modify transparent myCache local-zones { { name lz.example.net records add { "lz.example.net 60 IN A
127.0.0.1" "www.lz.example.net 300 IN A 127.0.0.2" } } }
```

Modifies DNS cache myCache by adding a local-zone lz.example.net with 2 resource records.

OPTIONS

answer-default-zones

Specifies whether the resolver cache answers queries for default zones: localhost, reverse 127.0.0.1 and ::1, and AS112 zones. The default value is no.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

local-zones

Zones and associated resource records for which the cache will provide Authoritative responses. Default is empty. This is intended for small, simple authoritative data configurations.

The local-zone name must be fully qualified and should be the apex of the zone. The local-zone type may be one of the following: deny, refuse, static, transparent, type-transparent, or redirect. Zero or more

resource records must be fully specified: name, ttl, class, type, and record data, separated by spaces, and within double quotes. For example, "www.example.net. 300 IN A 1.2.3.4".

For all local-zones types, if the DNS query matches, it is answered Authoritatively. How a non-matching query is handled depends on the local-zone type.

deny drops the query.

refuse sends a REFUSED response.

static sends either a NoData or NXDOMAIN response (includes SOA if present in local-zone).

transparent performs regular cache operation (i.e. transparent pass-through or iterative resolution) except for those query names which would result in NoData. This is the default local-zone type.

type-transparent Same as transparent but does not return NoData.

redirect returns responses with zone suffix record(s) for queries beneath that suffix. For example, a local-zone for example.com and a single A record for that name; queries for www.example.com or abc.www.example.com would return the single A record (both have the same suffix).

msg-cache-size
Specifies the maximum size in bytes of the DNS message cache. The default value is 1048576.

The BIG-IP system caches the messages in a DNS response in the message cache. After the maximum size of the cache is reached, when new or refreshed content is added to the cache, the expired and older content is removed from the cache. A higher maximum size allows more DNS responses to be cached and increases the cache hit percentage. A lower maximum size forces earlier eviction of cached content, but can lower the cache hit percentage.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

response-policy-zones
Adds, deletes or modifies the response policy zone to be used by this DNS Cache. Only a DNS Express zone configured as a response policy zone can be added.

The query name of a recursive DNS request without DNSSEC enabled is queried against the data in the response policy zone. If a match is found, the configured response policy action is taken.

action
The action to take upon a match. nxdomain results in an NXDOMAIN response given to the client. walled-garden results in a response with a CNAME to the walled-garden zone and an A or AAAA response matching the DNS query type. The default action is nxdomain.

walled-garden
A local zone configured in this cache that contains an A and/or AAAA record. This is typically used to redirect a user that requests resolution of a name contained in the RPZ database to a local server. This local server can display a message to the user and/or record the connection. Only A/AAAA/ANY requests are redirected, a request for any other type is answered with a NoData response. If a request is received for type A or AAAA but there are no records of that type configured, a NoData response is returned instead.

rrset-cache-size
Specifies the maximum size in bytes of the resource records set cache. The default value is 10485760.

The BIG-IP system caches the supporting records in a DNS response in the resource record cache. After the maximum size of the cache is reached, when new or refreshed content is added to the cache, the expired and older content is removed from the cache. A higher maximum size allows more DNS responses to be cached and increases the cache hit percentage. A lower maximum size forces earlier eviction of cached content, but can lower the cache hit percentage.

rrset-rotate
Specifies the resource record rotation method used within cached responses. The default value is none.

none Resource record order is not modified.

query-id Resource record order is a function of the client's query id.

SEE ALSO

create, delete, edit, glob, list, ltm dns cache resolver, ltm dns cache validating-resolver, show, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2016. All rights reserved.

Itm dns cache validating-resolver

NAME

validating-resolver - Configures a DNS cache with a resolver and validator on the BIG-IP(r) system.

MODULE

itm dns cache

SYNTAX

Configure the validating-resolver DNS cache component within the itm dns cache module using the syntax in the following sections.

CREATE/MODIFY

```
create validating-resolver [name]
modify validating-resolver [name]
options:
  allowed-query-time [integer]
  answer-default-zones [yes | no]
  app-service [[string] | none]
  description [[string] | none]
  dlvs-anchors {
    { [DNSKEY or DS RR string] ... }
  }
  forward-zones [add | delete | modify | replace-all-with] {
    [ [zone-name] ] {
options:
nameservers [add | delete | replace-all-with] {
  [ [IPv4address:port] | [IPv6address:port] ]
}
nameservers none
}
forward-zones none
ignore-cd [yes | no]
key-cache-size [integer]
local-zones [ [none] ]
[ { { name [dname] type [type] records [none | add { [RR string] ... } ] } ... } ] ]
max-concurrent-queries [integer]
max-concurrent-udp [integer]
max-concurrent-tcp [integer]
msg-cache-size [integer]
nameserver-cache-count [integer]
prefetch-key [yes | no]
randomize-query-name-case [yes | no]
response-policy-zones [add | delete | modify] {
  [zone-name] {
action [nxdomain | walled-garden]
walled-garden [local-zone]
}
}
response-policy-zones none
root-hints {
  { [IP address] ... }
}
route-domain [name]
rrset-cache-size [integer]
rrset-rotate [none | query-id]
trust-anchors {
  { [DNSKEY or DS RR string] ... }
}
unwanted-query-reply-threshold [integer]
use-ipv4 [yes | no]
use-ipv6 [yes | no]
use-tcp [yes | no]
use-udp [yes | no]
```

DISPLAY

```
list validating-resolver
list validating-resolver [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
show validating-resolver [name]
```

DELETE

```
delete validating-resolver [name]
```

DESCRIPTION

You can use the validating-resolver component to configure and view information about a validating recursive-

resolving DNS cache. A resolving and validating cache performs recursive resolution to fill its cache and uses DNSSEC to ensure the integrity of the data.

Important: When sizing caches, consider the total amount of memory available and how you wish to allocate memory for DNS caching. Note that cache sizing values are per-TMM process; therefore, a platform with eight TMMs consumes the amount of memory set for the resource record set cache times eight.

EXAMPLES

```
list validating-resolver myCache
```

Displays the properties of the validating recursive-resolving DNS cache myCache.

```
modify validating-resolver myCache local-zones { { name lz.example.net records add { "lz.example.net 60 IN A 127.0.0.1" "www.lz.example.net 300 IN A 127.0.0.2" } } }
```

Modifies DNS cache myCache by adding a local-zone lz.example.net with 2 resource records.

OPTIONS

allowed-query-time

Specifies the time allowed for a query to stay in the queue before it is replaced by a new query when the number of concurrent distinct queries exceeds the limit. The default value is 200 milliseconds.

answer-default-zones

Specifies whether the validating resolver cache answers queries for default zones: localhost, reverse 127.0.0.1 and ::1, and AS112 zones. The default value is no.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

dlv-anchors

Specifies the DNSKEY or DS resource records the BIG-IP system uses to establish DNSSEC trust with a DLV registry. The resource records must be specified in string format, for example, dig or drill format. The default is none.

forward-zones

Adds, deletes, modifies, or replaces a set of forward zones on a DNS Cache, by specifying zone name(s). A given zone name should only use the symbols allowed for a fully qualified domain name (FQDN), namely ASCII letters a through z, digits 0 through 9, hyphen -, and period .. For example site.example.com would be a valid zone name.

A DNS Cache configured with a forward zone will forward any queries that result in a cache-miss (the answer was not available in the cache) and match a configured zone name, to the nameserver specified on the zone. If no nameservers are specified on the zone, an automatic SERVFAIL is returned. When a forward zone's nameserver returns a valid response to the DNS Cache, that response is cached and then returned to the requester.

nameservers

Adds, deletes, or replaces a set of nameservers in a forward zone on a DNS Cache. A nameserver is represented by an IP address and port in the format [IPv4:port] or [IPv6.port], for example 10.10.10.10:53 or 2001::1:ff:53, respectively.

If more than one nameserver is listed for a given forward zone, a matching query will be sent to the nameserver that is currently deemed the most responsive (based on RTTs). If no response is received within a certain window of time, the DNS Cache will resend the query to another nameserver with an increased wait window until a response is received.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ignore-cd

When enabled, the system ignores the Checking Disabled setting on client queries, performs validation, and returns only secure answers. The default value is no.

key-cache-size

Specifies the maximum size in bytes of the DNSKEY cache. The default value is 1048576.

local-zones

Zones and associated resource records for which the cache will provide Authoritative responses. Default is empty. This is intended for small, simple authoritative data configurations.

The local-zone name must be fully qualified and should be the apex of the zone. The local-zone type may be one of the following: deny, refuse, static, transparent, type-transparent, or redirect. Zero or more resource records must be fully specified: name, ttl, class, type, and record data, separated by spaces, and within double quotes. For example, "www.example.net. 300 IN A 1.2.3.4".

For all local-zones types, if the DNS query matches, it is answered Authoritatively. How a non-matching query is handled depends on the local-zone type.

deny drops the query.

refuse sends a REFUSED response.

static sends either a NoData or NXDOMAIN response (includes SOA if present in local-zone).

transparent performs regular cache operation (i.e. transparent pass-through or iterative resolution) except for those query names which would result in NoData. This is the default local-zone type.

type-transparent Same as transparent but does not return NoData.

redirect returns responses with zone suffix record(s) for queries beneath that suffix. For example, a local-zone for example.com and a single A record for that name; queries for www.example.com or abc.www.example.com would return the single A record (both have the same suffix).

max-concurrent-queries

Specifies the maximum number of concurrent distinct queries used by the resolver. A query is identified by query name, type and class. If the number of distinct queries exceeds this limit, the resolver replaces the earliest query in the queue with the new query if it has been in the queue longer than the allowed time. The default value is 1024.

max-concurrent-tcp

Specifies the maximum number of concurrent TCP flows used by the resolver. The default value is 20.

max-concurrent-udp

Specifies the maximum number of concurrent UDP flows used by the resolver. The default value is 8192.

msg-cache-size

Specifies the maximum size in bytes of the DNS message cache. The default value is 1048576.

The BIG-IP system caches the messages in a DNS response in the message cache. After the maximum size of the cache is reached, when new or refreshed content is added to the cache, the expired and older content is removed from the cache. A higher maximum size allows more DNS responses to be cached and increases the cache hit percentage. A lower maximum size forces earlier eviction of cached content, but can lower the cache hit percentage.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

nameserver-cache-count

Specifies the maximum number of DNS nameservers for which the BIG-IP system caches connection and capability data. The default value is 16536 entries.

prefetch-key

When enabled, the validating resolver fetches the DNSKEY early in the validation process. Disable this setting when you want to reduce resolver traffic, but understand that a client may have to wait for the validating resolver to perform a key lookup. The default value is yes.

randomize-query-name-case

When enabled, the resolver randomizes the case of query names. The default value is yes.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

response-policy-zones

Adds, deletes or modifies the response policy zone to be used by this DNS Cache. Only a DNS Express zone configured as a response policy zone can be added.

The query name of a recursive DNS request without DNSSEC enabled is queried against the data in the response policy zone. If a match is found, the configured response policy action is taken.

action

The action to take upon a match. nxdomain results in an NXDOMAIN response given to the client. walled-garden results in a response with a CNAME to the walled-garden zone and an A or AAAA response matching the DNS query type. The default action is nxdomain.

walled-garden

A local zone configured in this cache that contains an A and/or AAAA record. This is typically used to redirect a user that requests resolution of a name contained in the RPZ database to a local server. This local server can display a message to the user and/or record the connection. Only A/AAAA/ANY requests are redirected, a request for any other type is answered with a NoData response. If a request is received for type A or AAAA but there are no records of that type configured, a NoData response is returned instead.

root-hints

Specifies the IP addresses of DNS servers that the BIG-IP system considers authoritative for the DNS root nameservers.

Important:By default, the BIG-IP system uses the DNS root nameservers published by InterNIC.

Caution:When you add DNS root nameservers, the BIG-IP system no longer uses the default nameservers published by InterNIC, but instead uses the nameservers you add as authoritative for the DNS root nameservers.

route-domain

Specifies the route domain the resolver uses for outbound traffic. The default value is the default route domain.

rrset-cache-size

Specifies the maximum size in bytes of the resource records set cache. The default value is 10485760.

The BIG-IP system caches the supporting records in a DNS response in the resource record cache. After the maximum size of the cache is reached, when new or refreshed content is added to the cache, the expired and older content is removed from the cache. A higher maximum size allows more DNS responses to be cached and increases the cache hit percentage. A lower maximum size forces earlier eviction of cached content, but can lower the cache hit percentage.

`rrset-rotate`

Specifies the resource record rotation method used within cached responses. The default value is none.

none Resource record order is not modified.

`query-id` Resource record order is a function of the client's query id.

`trust-anchors`

Specifies the DNSKEY or DS resource records the BIG-IP system uses to establish DNSSEC trust with a specific DNS zone. The resource records must be specified in string format, for example, dig or drill format. The default value is none.

`unwanted-query-reply-threshold`

The system always rejects unsolicited replies. The default value of 0 (off) indicates the system does not generate SNMP traps or log messages when rejecting unsolicited replies.

Change the default value to monitor for unsolicited DNS replies. This alerts you to a potential security attack, such as cache poisoning or DOS. For example, if you specify a value of 1,000,000, each time the system receives 1,000,000 unsolicited replies, it generates an SNMP trap and log message. The default value is 0 (off).

`use-ipv4`

When enabled, the resolver sends DNS queries to IPv4 addresses. The default value is yes.

`use-ipv6`

When enabled, the resolver sends DNS queries to IPv6 addresses. The default value is yes.

`use-tcp`

When enabled, the resolver can send queries over the TCP protocol. The default value is yes.

`use-udp`

When enabled, the resolver can send queries over the UDP protocol. The default value is yes.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm dns cache transparent`, `ltm dns cache resolver`, `show`, `modify`, `regex`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2016. All rights reserved.

BIG-IP 2017-09-07 ltm dns cache validating-resolver(1)

ltm dns dns-express-db

NAME

`dns-express-db` - Loads the DNS Express data file.

MODULE

ltm dns

SYNTAX

`load dns-express-db`

DESCRIPTION

The `dns-express-db` component within the ltm dns module is used to load the DNS Express data file `/var/db/tmmdns.bin`. The file is only loaded if it has been modified.

EXAMPLES

`load dns-express-db`

Loads the DNS Express file from disk into the running configuration.

SEE ALSO

`load`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 Itm dns dns-express-db(1)

Itm dns dnssec key

NAME

key - Configures DNSSEC keys on the BIG-IP(r) system.

MODULE

itm dns dnssec

SYNTAX

Configure the key component within the Itm dns dnssec module using the syntax in the following sections.

CREATE/MODIFY

create key [name]

modify key [name]

options:

algorithm [rsasha1 | rsasha256 | rsasha512]

app-service [[string] | none]

bitwidth [512 | 1024 | 2048 | 4096]

certificate-file [string]

description [string]

[enabled | disabled]

expiration-period [integer]

generation {

[[generation-id]] }

options:

expiration [date:time]

rollover [date:time]

key-file [string]

key-type [ksk | zsk]

rollover-period [integer]

signature-pub-period [integer]

signature-valid-period [integer]

ttl [integer]

use-fips [external | internal | none]

edit key [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list key

list key [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete key [name]

DESCRIPTION

You can use the key component to configure DNSSEC zone signing and key signing keys, and to view information about the keys.

EXAMPLES

create key ksk1

Creates the key signing key, ksk1, using the system default values.

create key zsk1

Creates the zone signing key, zsk1, using the system default values.

list key my_key

Displays the properties of the DNS security key my_key.

OPTIONS

algorithm

Specifies the algorithm to use to generate the key. The default value is RSASHA1.

app-service

Specifies the name of the application service to which the key belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the key. Only the application service can modify or delete the key.

bitwidth

Specifies the length of the key you want to generate. The default value is 1024. If a key is manually managed, MCPD will derive this value from the file and override any user defined value.

certificate-file

Specifies the file containing the public key. Fields certificate-file and key-file are required for manual DNSSEC key import.

description

User defined description.

[enabled | disabled]

Specifies whether the key is enabled or disabled.

expiration-period

Specifies the life of the key in d:h:m:s, h:m:s, m:s, or seconds. At the end of the period, the system deletes the expired generation of the key. This value must be greater than the value of the rollover-period option. The difference between the two periods must be more than the value of the ttl option.

The default value is 0 (zero), which indicates unset, and thus the key does not expire.

generation

Displays the generation of the key, including the following:

creator

Hostname of BIG-IP system that created this generation.

expiration

The date and time that this generation of the key expires. This can be modified and is in the following format: yyyy-mm-dd:hh:mm:ss.

handle

The handle of a generation of a key that is used for internal interactions with the key subsystem (for example, HSM for FIPS).

key-tag

The hash identifier of the DNSKEY. This can be used to identify which DNSKEY was used to generate a given RRSIG.

pub-text

The text of the public portion of the DNSSEC Key Generation.

rollover

The date and time that the generation of the key rolls over to a new key. This can be modified and is in the following format: yyyy-mm-dd:hh:mm:ss.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

key-file

Specifies the file containing the private key. Fields certificate-file and key-file are required for manual DNSSEC key import.

key-type

Specifies whether the key is of type ksk or zsk. The default value is zsk.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rollover-period

Specifies the amount of time, in d:h:m:s, h:m:s, m:s, or seconds, before the system generates another generation of the key. At the end of the period, the system creates a new generation of the key. Two generations of the key exist during the time between the end of the rollover period and the end of the expiration period.

This value must be greater than or equal to one third of the value of the expiration-period option, and less than the value of the expiration period option. The difference between the two periods must be more than the value of the ttl option.

The default value is 0 (zero), which indicates unset, and thus the key does not roll over.

signature-pub-period

Specifies the amount of time, in d:h:m:s, h:m:s, m:s, or seconds, before the system publishes another generation of the signature. At the end of the period, the system creates a new signature.

This value must be less than the value of the signature-valid-period option. The default value is 403200 seconds.

signature-valid-period

Specifies the amount of time, in d:h:m:s, h:m:s, m:s, or seconds, that the signature is valid. The validity period will begin when the signature is generated but the inception time of the signature will be back-dated by one hour, to allow for clock skew on the validator. At the end of the period, the Global Traffic Manager no longer uses the expired signature. The default value is 604800 seconds.

ttl Specifies the amount of time, in d:h:m:s, h:m:s, m:s, or seconds, that a DNS server can cache the key. The default value is 86400.

The value of the **ttl** option must be less than the difference between the values of the **rollover-period** and **expiration-period** options.

0 seconds indicates that the key is not cached.

use-fips

Specifies the type of FIPS-compliant hardware security module to use when storing, and signing with, the private key. The default value is none. The choice of external attempts to use a network-attached FIPS device if configured; otherwise internal uses the FIPS device within the BIG-IP.

If this option is set to **internal** or **external** and a FIPS device is not present, the system automatically resets the value to **none**.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2019-05-13 itm dns dnssec key(1)

Itm dns dnssec zone

NAME

`zone` - Configures DNSSEC zones on the BIG-IP(r) system.

MODULE

`itm dns dnssec`

SYNTAX

Configure the zone component within the `itm dns dnssec` module using the syntax in the following sections.

CREATE/MODIFY

```
create zone [name]
modify zone [name]
options:
  app-service [[string] | none]
  description [string]
  [enabled | disabled]
  ds-algorithm [ sha1 | sha256 ] DEPRECATED - see ds-algorithms
  ds-algorithms [ add | delete | replace-all-with ] {
    [ sha1 | sha256 ] ...
  }
  external-delegations
    [add | delete | modify | replace-all-with] {
[DNS zone name] {
  options:
    ds-records
      [add | delete | modify | replace-all-with] {
[ DS record ] ...
      }
    secure [ enabled | disabled ]
}
}
  indicate-authenticated [ enabled | disabled ]
  keys
    [add | delete | modify | replace-all-with] {
[key name ...]
  }
  keys none
  nsec3-algorithm [ SHA1 ]
  nsec3-iterations [unsigned integer]
  publish-cds-cdnskey [ enabled | disabled ]

edit zone [ [ [name] | [glob] | [regex] ] ... ]
options:
```

all-properties
non-default-properties

reset-stats zone
reset-stats zone [[[name] | [glob] | [regex]] ...]

DISPLAY
list zone
list zone [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
seps
show zone [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
global
field-fmt

DELETE
delete zone [name]

DESCRIPTION
You can use the zone component to configure and view information about a DNSSEC zone.

EXAMPLES
list zone mySecureZone

Displays the properties of the DNSSEC zone named mySecureZone.

OPTIONS
app-service
Specifies the name of the application service to which the zone belongs. The default value is none. Note:
If the strict-updates option is enabled on the application service that owns the object, you cannot
modify or delete the zone. Only the application service can modify or delete the zone.

description
User defined description.

ds-algorithm
This option is deprecated in v14.0.0 and is replaced by ds-algorithms. Specifies the hash algorithm to
use when creating the Delegation Signer (DS) resource record. The default value is sha1.

ds-algorithms
Specifies the hash algorithms to use when creating Delegation Signer (DS) resource records. The default
value is sha1. A DS record is generated in a given SEP for each algorithm that is configured.

[enabled | disabled]
Specifies whether the DNSSEC zone is enabled or disabled.

Note: You must associate both a key signing and a zone signing key with the zone before complete signing
of client requests can occur.

external-delegations
Specifies the names of delegated subzones of this zone, where the BIG-IP is not responsible for the
DNSSEC signing.

ds-records
Specifies the DNSSEC delegation signer (DS) resource records (RRs) that correspond to the Key-
Signing-Keys (KSKs) of the external delegated zone. They indicate that the external delegated zone
is DNSSEC enabled. These records are used to establish the DNSSEC chain of trust from zone to
subzone.

secure
Specifies whether or not the external delegation is secured through the use of DS records. Default
value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression
syntax.

indicate-authenticated
The Authenticated Data (AD) flag is TRUE for DNSSEC zone authoritative answers when this setting is
enabled. The default value is disabled.

keys Specifies the keys that you want to configure for the zone.

name Specifies a unique name for the component. This option is required for the commands create, delete, and
modify.

nsec3-algorithm
Specifies the hash algorithm to use when creating the Next Secure (NSEC3) resource record. The default
value is SHA1. Other algorithms are not currently supported, so selecting SHA256 will revert to SHA1 with
a warning message.

nsec3-iterations
Specifies the number of times to hash the Next Secure (NSEC3) names. The default value is 1.

`publish-cds-cdnskey`

Specifies whether or not we will respond to CDS and CDNSKEY type queries for the DNSSEC Zone. The default value is disabled.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`secure-delegations`

Specifies the DNSSEC zones on the BIG-IP that are delegated subzones of the zone as determined by the name of the zones. This list is read-only and automatically generated based on the DNSSEC Zones configured on the BIG-IP.

`seps` Displays the Secure Entry Point(s) (DS and DNSKEY resource records used as client trust anchors) of the zone. This list is read-only and automatically generated based on the DNSSEC Key Key-Signing-Keys (KSKs) configured on a DNSSEC Zone.

Each list entry includes the following attributes:

`dnskey`

String representation of the DNSKEY resource record. Note this will be a Key-Signing-Key (KSK).

`ds` This option is deprecated in v14.0.0 and is replaced by `ds-records`. String representation of the DS resource record.

`ds-records`

String representations of DS resource records. There will be one DS record for each `ds-algorithm` configured on the DNSSEC Zone.

`generation-id`

Generation ID of DNSSEC Key used to create the SEP.

`key-name`

Name of DNSSEC Key which was used to create the SEP.

`xfr-primary-soa-serial`

The learned zone SOA serial number from the primary server.

`xfr-soa-serial`

The advertised zone SOA serial number to all clients.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015-2016. All rights reserved.

BIG-IP 2018-09-12 ltm dns dnssec zone(1)

ltm dns nameserver

NAME

nameserver - Configures DNS nameservers on the BIG-IP(r) system.

MODULE

ltm dns

SYNTAX

Configure the nameserver component within the ltm dns module using the syntax in the following sections.

CREATE/MODIFY

`create nameserver [name]`

`modify nameserver [name]`

options:

`address [ip address]`

`app-service [[string] | none]`

`description [[string] | none]`

`port [unsigned integer]`

`route-domain [route-domain name | none]`

`tsig-key [tsig-key name | none]`

`edit nameserver [[[name] | [glob] | [regex]] ...]`

options:

all-properties
non-default-properties

reset-stats nameserver
reset-stats nameserver [[[name] | [glob] | [regex]] ...]

DISPLAY
list nameserver
list nameserver [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
show nameserver
show nameserver [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
global
field-fmt

DELETE
delete nameserver [name]

DESCRIPTION

You can use the nameserver component to configure nameservers and to view information about the nameservers.

EXAMPLES

```
create nameserver myNameserver address 127.0.0.1 port 53
```

Creates the nameserver, myNameserver, given the address and port.

```
list nameserver myNameserver
```

Displays the properties of the nameserver myNameserver.

OPTIONS

address
Specifies the IP address of the nameserver. The default value is 127.0.0.1.

app-service
Specifies the name of the application service to which the nameserver belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the nameserver. Only the application service can modify or delete the nameserver.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

port Specifies the service port of the nameserver. The default value is 53.

route-domain
Specifies the route domain that the nameserver uses for outbound traffic. The default value is the default route domain.

tsig-key
Specifies the TSIG key used to communicate with this nameserver for zone transfers. If the nameserver is a client, then this TSIG key is used to verify the query and sign the response. If the nameserver is a transfer target for DNS Express nameserver, then this TSIG key should match that of the master nameserver.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2017-09-07 Itm dns nameserver(1)

ltm dns tsig-key

NAME

tsig-key - Configures TSIG keys on the BIG-IP(r) system.

MODULE

ltm dns

SYNTAX

Configure the tsig-key component within the ltm dns module using the syntax in the following sections.

CREATE/MODIFY

create tsig-key [name]

modify tsig-key [name]

options:

algorithm [hmacmd5 | hmacsha1 | hmacsha256]

app-service [[string] | none]

description [[string] | none]

secret [string]

edit tsig-key [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tsig-key

list tsig-key [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete tsig-key [name]

DESCRIPTION

You can use the tsig-key component to configure TSIG keys and to view information about the keys.

EXAMPLES

```
create tsig-key myKey algorithm hmacmd5 secret ABCDEFG
```

Creates the TSIG key, myKey, given the algorithm and secret (both required).

```
list tsig-key myKey
```

Displays the properties of the TSIG key myKey.

OPTIONS

algorithm

Specifies the algorithm to use to generate the key.

app-service

Specifies the name of the application service to which the TSIG key belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the TSIG key. Only the application service can modify or delete the TSIG key.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Specifies the string representation of the key's shared secret.

SEE ALSO

create, delete, edit, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

Itm dns zone

NAME

zone - Configures zones on the BIG-IP(r) system.

MODULE

itm dns

SYNTAX

Configure the zone component within the Itm dns module using the syntax in the following sections.

CREATE/MODIFY

create zone [name]

modify zone [name]

options:

app-service [[string] | none]

description [[string] | none]

dns-express-allow-notify [add | delete | none | replace-all-with] {
[IP Address]

}

dns-express-enabled [yes | no]

dns-express-notify-action [consume | bypass | repeat]

dns-express-notify-tsig-verify [yes | no]

dns-express-server [server name | none]

response-policy [yes | no]

server-tsig-key [tsig-key name | none]

transfer-clients [add | delete | none | replace-all-with] {
[server name]

}

edit zone [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats zone

reset-stats zone [[[name] | [glob] | [regex]] ...]

DISPLAY

list zone

list zone [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show zone [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete zone [name]

DESCRIPTION

You can use the zone component to configure and view information about a zone.

EXAMPLES

```
list zone myZone
```

Displays the properties of the zone named myZone.

```
create zone myZone transfer-clients add { nameserver1 nameserver2 }
```

Creates a zone named myZone, which allows zone data to be transferred to nameserver1 and nameserver2.

OPTIONS

app-service

Specifies the name of the application service to which the zone belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the zone. Only the application service can modify or delete the zone.

description

User defined description.

dns-express-allow-notify

Specifies a list of IP addresses, in addition to the DNS Zone's DNS-Express Server address, which are

allowed to notify the BIGIP of DNS Zone changes. A notify message coming from an IP which is neither the address of the zone's DNS Express server nor an address in this list will be dropped by the BIGIP.

`dns-express-enabled` [yes | no]

Specifies whether DNS Express is enabled to process queries for this zone. The default value is yes.

`dns-express-notify-action` [consume | bypass | repeat]

Action to take when a NOTIFY query is received for a configured zone. Options are consume, bypass, and repeat. Default is consume, meaning the NOTIFY query is seen only by DNS Express. bypass means the query will NOT go to DNS Express, but any backend DNS resource (subject to DNS profile unhandled-query-action). repeat means the NOTIFY will go to both DNS Express and any backend DNS resource. If TSIG is configured, the signature is only validated for consume and repeat actions. NOTIFY responses are assumed to be sent by the backend DNS resource, except when the action is consume and DNS Express will generate a response.

`dns-express-notify-tsig-verify`

Verify NOTIFY query TSIG for a DNS Express zone. Default is yes.

`dns-express-server`

Specifies the server from which to retrieve zone information for DNS Express.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

Note: A successful zone transfer must occur before this zone can service DNS requests.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`response-policy`

Specifies if this is a response policy zone. If this is set to yes, this zone may be assigned as an RPZ to a DNS Cache. Default is no.

`server-tsig-key`

Specifies the server side TSIG key associated with the DNS zone. It should match the TSIG Key associated with the master name servers.

`transfer-clients`

Specifies the nameservers allowed to transfer the zone from BIGIP.

SEE ALSO

create, delete, edit, glob, list, show, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014. All rights reserved.

BIG-IP 2017-09-07 ltm dns zone(1)

ltm eviction-policy

NAME

eviction-policy - Configures eviction policies to determine when and how to terminate connections.

MODULE

ltm

SYNTAX

Configure the eviction-policy component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

create eviction-policy [name]

modify eviction-policy [name]

options:

description [string]

app-service [[string] | none]

low-water [integer]

high-water [integer]

slow-flow {

enabled [true | false]

eviction-type [count | percent]

```

    grace-period [integer]
    maximum [integer]
    threshold-bps [integer]
    throttling [enabled | disabled]
  }
  strategies {
    bias-bytes {
  delay [integer]
  enabled [true | false]
    }
    bias-idle {
  enabled [true | false]
    }
    bias-oldest {
  enabled [true | false]
    }
    low-priority-geographies {
  countries [add | delete | modify | replace-all-with] {
    [country-code] ...
  }
  enabled [true | false]
    }
    low-priority-port {
  enabled [true | false]
  ports [add | delete | modify | replace-all-with] {
    [ [name] ] {
      app-service [[string] | none]
      port-number [name | integer]
      protocol [any | sctp | tcp | udp]
    } ...
  }
    }
    low-priority-route-domain {
  enabled [true | false]
  names [add | delete | modify | replace-all-with] {
    [ [route domain name] ] ...
  }
    }
    low-priority-virtual-server {
  enabled [true | false]
  names [add | delete | modify | replace-all-with] {
    [ [virtual server name] ] ...
  }
    }
  }
}

```

DISPLAY

```
list eviction-policy
```

```
list eviction-policy [ [ [name] | [glob] | [regex] ] ...]
```

```
options:
```

```
  all-properties
  partition
```

```
show eviction-policy
```

```
show eviction-policy [ [ [name] | [glob] | [regex] ] ...]
```

```
options:
```

```
  all-properties
  default
```

DELETE

```
delete eviction-policy [name]
```

DESCRIPTION

You use the eviction policy to specify which flows to terminate when the connection limits for the box are approached. The eviction policy contains strategies which select the flows to terminate. Additionally, the eviction policy defines parameters used to determine when flows are considered to be slow. Slow flows are terminated according to the policy, even when the Big-IP is not under duress and the connection limits are not approached.

When applied to the global context, the eviction policy low-water and high-water limits are with respect to memory available on the Big-IP. When applied to a virtual server or a route domain, the limits are with respect to the connection limit on the virtual server or route domain, respectively.

It is possible to monitor slow flows, accumulating metrics on the number of flows under the designated slow flow transfer limit. To do so, enable slow-flows, but disable throttling on the slow flow.

Note: Monitoring or killing slow flows will incur a performance penalty.

Note: The strategies applied here work on a cyclic sweep of all connections on the Big-IP. These do not run at the granularity to guarantee eviction of a particular flow or type of flow, but are statistical and opportunistic.

The bias-bytes algorithm attempts to select the flows that have sent and received the fewest bytes on the connection.

The bias-idle algorithm attempts to select the flows that have been idle the longest.

The bias-oldest algorithm attempts to select the oldest flows.

The low-priority-geographies algorithm selects flows that are in low-priority geographies according to the GeolP database loaded onto the box.

The low-priority-port algorithm selects flows that are in the provided list of low-priority ports and protocols.

The low-priority-route-domain algorithm selects flows that are in the provided list of low-priority route domains.

The low-priority-virtual-server algorithm selects flows that are in the provided list of low-priority virtual servers.

EXAMPLES

```
create eviction-policy my_eviction_policy { low-water 70 high-water 80 slow-flow { enabled true threshold-bps 50 throttling disabled } strategies { bias-idle { enabled true } } }
```

Creates an eviction policy named `my_eviction_policy`, which accumulates statistics on the current number of slow flows but does not terminate any flows that are considered slow. The `bias-idle` algorithm is used to kill flows when the limits on the context are approached. The aggressive sweeper will activate at 80 percent of capacity, and deactivate when load is reduced to 70 percent of capacity.

```
modify eviction-policy my_eviction_policy { strategies { low-priority-geographies { enabled true countries replace-all-with { AZ BZ } } } }
```

Modifies the eviction policy named `my_eviction_policy`, enabling the low-priority geography strategy and dropping flows from Azerbaijan (AZ) and Belize (BZ).

```
delete eviction-policy my_eviction_policy
```

Deletes the eviction policy named `my_eviction_policy`.

OPTIONS

`description`

Provides a user-defined description for the policy.

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`low-water`

Specifies the aggressive sweeper activation threshold as a percentage of total capacity. The allowable range is 50 - 100, and the `low-water` value must be lower than or equal to the `high-water` value. To disable killing flows when limits are met or exceeded, set the `low-water` and `high-water` to 100 percent. Default value is 85 percent.

`high-water`

Specifies the target maximum load on the context. The adaptive reaper will be more aggressive as this limit is approached. The allowable range is 50 - 100, and the `high-water` value must be higher than or equal to the `low-water` value. To disable killing flows when limits are met or exceeded, set the `low-water` and `high-water` to 100 percent. Defaults to 95 percent.

`slow-flow`

Specifies whether to monitor and possibly remove flows considered to be slow.

`enabled`

If true, the `slow-flow` monitoring and possible removal are activated. If false, the remaining `slow-flow` attributes are unused. The default value is false.

`eviction-type`

Indicates whether the threshold is based on an absolute count of slow flows, or a percentage of the total flows on the context where the eviction policy is applied. There is no default value.

`grace-period`

Specifies the minimum age of a slow flow before the flow is killed.

`maximum`

Provides the count or percentage at which slow flows will be killed. If `eviction-type` is `count`, this value is the absolute number of slow flows allowed; if `eviction-type` is `percent`, this is the percentage of flows on the context that are allowed to fall under the `threshold-bps` before being killed.

`threshold-bps`

Provides the threshold under which flows are considered to be slow, in bytes per second.

`throttling`

Indicates whether to kill flows that are considered slow. If set to `enabled`, flows that fall under the `threshold-bps` are subject to being killed according to the defined maximum number of flows.

`strategies`

Defines the strategies to be used to select flows for eviction in the eviction policy.

`bias-bytes`

Defines how to use the `bias-bytes` eviction strategy.

delay

The delay allowed for new flows to transfer content, to prevent killing infant flows.

enabled

Specifies whether to use the bias-bytes algorithm. If false, the bias-bytes algorithm is not used. The default value is false.

bias-idle

Defines how to use the bias-idle eviction strategy.

enabled

Specifies whether to use the bias-idle algorithm. If false, the bias-idle algorithm is not used. The default value is false.

bias-oldest

Defines how to use the bias-oldest eviction strategy.

enabled

Specifies whether to use the bias-oldest algorithm. If false, the bias-oldest algorithm is not used. The default value is false.

low-priority-geographies

Defines how to use the low-priority geographies eviction strategy.

countries

Provides a list of country codes considered low-priority candidates to evict, based on GeoIP information.

enabled

Specifies whether to use the low-priority-geographies algorithm. If false, the low-priority-geographies algorithm is not used. The default value is false.

low-priority-port

Defines how to use the low-priority-port eviction strategy.

enabled

Specifies whether to use the low-priority-port algorithm. If false, the low-priority-port algorithm is not used. The default value is false.

ports

Provides a list of ports considered low-priority candidates to evict.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

port-number

Specifies the port number considered low-priority. The number provided can be either a number (e.g., 80) or a name (e.g., http).

protocol

Specifies the protocol considered low-priority. The default value is any.

low-priority-route-domain

Defines how to use the low-priority-route-domain eviction strategy.

enabled

Specifies whether to use the low-priority-route-domain algorithm. If false, the low-priority-route-domain algorithm is not used. The default value is false.

names

Specifies a list of route domain names considered to be low-priority candidates to evict.

low-priority-virtual-server

Defines how to use the low-priority-virtual-server eviction strategy.

enabled

Specifies whether to use the low-priority-virtual-server algorithm. If false, the low-priority-virtual-server algorithm is not used. The default value is false.

names

Specifies a list of virtual server names considered to be low-priority candidates to evict.

SEE ALSO

create, delete, edit, list, modify, show, tmsh, regex, sys geoip

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2015. All rights reserved.

Itm global-settings connection

NAME

connection - Configures the global settings that pertain to connections for the BIG-IP(r) and VIPRION(r) local traffic management systems.

MODULE

itm global-settings

SYNTAX

Configure the connection component within the Itm global-settings module using the syntax shown in the following sections.

MODIFY

modify connection

options:

adaptive-reaper-hiwater [integer]
adaptive-reaper-lowwater [integer]
auto-last-hop [disabled | enabled]
default-vs-syn-challenge-threshold [infinite | integer]
global-flow-eviction-policy [name]
global-syn-challenge-threshold [infinite | integer]
syncookies-threshold [integer]
vlan-keyed-conn [disabled | enabled]
vlan-syn-cookie [disabled | enabled]

DISPLAY

list connection

list connection [option name]

show running-config connection

show running-config connection [option name]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the connection component to modify how the system processes connections.

EXAMPLES

modify connection auto-last-hop disabled

Specifies that the system does not automatically map the last hop for pools.

list connection

Displays the global settings for how the system processes connections.

OPTIONS

adaptive-reaper-hiwater

IMPORTANT This command has been deprecated (as of 11.6.0). Please use Itm eviction-policy instead. Specifies, in a percentage, the memory usage at which the system stops establishing new connections. Once the system meets the reaper high-water mark, the system does not establish new connections until the memory usage drops below the reaper low-water mark. The adaptive reaper settings help mitigate the effects of a denial-of-service attack.

The available range is 85 - 100. The default value is 95. To disable the adaptive reaper, set the high-water mark to 100.

adaptive-reaper-lowwater

IMPORTANT This command has been deprecated (as of 11.6.0). Please use Itm eviction-policy instead. Specifies, in percent, the memory usage at which the system silently purges stale connections, without sending reset packets (RST) to the client. If the memory usage remains above the low-water mark after the purge, then the system starts purging established connections closest to their service timeout.

The available range is 70 - 100. The default value is 85. To disable the adaptive reaper, set the low-water mark to 100.

auto-last-hop

Specifies that the system automatically maps the last hop for pools. The default value is enabled.

default-vs-syn-challenge-threshold

Specifies the default value of per-virtual server SYN Cookie activation threshold per chassis. The default value is infinite. The valid range is 128 - 1024K or infinite (encoded as 0).

global-flow-eviction-policy

Specifies the flow eviction policy to use when approaching memory usage limits. The settings in the policy determine the adaptive reaper high and low water marks, and help determine which client connections to terminate when memory limits have exceeded the "low-water" threshold in the eviction policy. The settings help mitigate the effects of a denial-of-service attack.

global-syn-challenge-threshold

Specifies the default value of the global SYN Cookie activation threshold per TMM. The default value is 64K. The valid range is 2048 - 4096K or infinite (encoded as 0).

syncookies-threshold

This option is deprecated in version 13.0.0 and is replaced by default-vs-syn-challenge-threshold. Specifies the number of new or untrusted TCP connections that can be established before the system activates the SYN Cookies authentication method for subsequent TCP connections. The default value is 16384.

vlan-keyed-conn

Enables or disables VLAN-keyed connections. You use VLAN-keyed connections when traffic for the same connection must pass through the system several times, on multiple pairs of VLANs (or in different VLAN groups). The default value is enabled.

vlan-syn-cookie

Enables or disables the hardware per-VLAN SYN cookie protection on platforms with supported firmware. The default value is enabled.

SEE ALSO

list, ltm node, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013, 2016. All rights reserved.

BIG-IP 2016-09-06 ltm global-settings connection(1)

ltm global-settings general

NAME

general - Configures the general properties for the BIG-IP(r) and VIPRION(r) local traffic management systems.

MODULE

ltm global-settings

SYNTAX

Configure the general component within the ltm global-settings module using the syntax shown in the following sections.

MODIFY

modify general

options:

gratuitous-arp-rate [integer value: 0 ~ 2147483647]
l2-cache-timeout [integer value: 0 ~ 2147483647]
maintenance-mode [disabled | enabled]
mgmt-auto-lasthop [disabled | enabled]
share-single-mac [unique | global | vmw-compat]
snat-packet-forward [disabled | enabled]

DISPLAY

list general

list general [option name]

show running-config general

show running-config general [option name]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the general component to modify how the system processes local traffic.

EXAMPLES

modify general maintenance-mode enabled

Places the Local Traffic Manager system in maintenance mode.

list general

Displays the general properties of the local traffic management system.

OPTIONS

gratuitous-arp-rate

Specifies how fast gratuitous ARPs can be sent. If it is 0, then gratuitous ARPs are sent without pause. Otherwise, it specifies how many gratuitous ARPs can be sent every second. The default value is 0. The range is 0 (zero) to 2147483647."

l2-cache-timeout

Specifies, in seconds, the amount of time that records remain in the Layer 2 forwarding table, when the MAC address of the record is no longer detected on the network.

The default value is 300 seconds. The range is 0 (zero) to 2147483647 seconds.

maintenance-mode

Specifies, when enabled, that the unit is in maintenance mode. In maintenance mode, the system stops accepting new connections and slowly finishes off existing connections.

The default value is disabled.

mgmt-auto-lasthop

Specifies, when enabled, that auto-lasthop should be used for incoming traffic to the management port. That means return traffic will be sent to the MAC address that originated the traffic rather than looked up using the routing and ARP/NDP tables.

The default value is enabled.

share-single-mac

Specifies the Media Access Control address (MAC address) that the system assigns to a VLAN. The default value is unique, which indicates that a VLAN uses a unique MAC address from the pool of mac addresses assigned to each hardware platform. The global value indicates that all of the VLANs on the system use the same MAC address. The vmw-compat value indicates that the MAC address of a vlan is allocated in a manner compatible with VMware(tm) vSwitch, and restricts vlans to a single interface, with no trunks allowed. Changing the value of this feature requires a manual restart of all TMOS daemons.

snat-packet-forward

Enables or disables SNAT packet forwarding. The default value is enabled.

SEE ALSO

list, ltm node, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013, 2016. All rights reserved.

BIG-IP 2017-06-28 ltm global-settings general(1)

ltm global-settings rule

NAME

rule - Configures the iRule properties for the BIG-IP(r) and VIPRION(r) local traffic management systems.

MODULE

ltm global-settings

SYNTAX

Configure the rule component within the ltm global-settings module using the syntax shown in the following sections.

MODIFY

modify rule

options:

rule-aborted-log-ratio [integer value: 0 ~ 2147483647]

DISPLAY

list rule

list rule [option name]

show running-config rule

show running-config rule [option name]

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the rule component to modify how the system processes iRule conditions.

EXAMPLES

modify rule rule-aborted-log-ratio 5

Instructs the system to generate one log message for every 5 aborted iRule executions.

list rule

Displays the iRule properties of the local traffic management system.

OPTIONS

rule-aborted-log-ratio

Specifies the ratio of log messages generated when iRule executions are aborted. If set to 1, every aborted execution is logged. If set higher, every n aborted executions will result in one log message. If set to 0, no message is generated when executions are aborted.

Note, this setting is per TMM across all iRules; the implication is that if the value is set to a number greater than 1 in order to reduce the rate of log messages, all aborted rule executions are affected.

If there are multiple rules that could cause this condition where one rule is causing an excessive amount of such logs (and this condition is considered non important), and another rule where this condition signifies a problem, and the ratio is set to a large number, the occurrence of this condition for the second rule could be lost.

The default value is 1. The range is 0 (zero) to 2147483647.

SEE ALSO

list, ltm node, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2017. All rights reserved.

BIG-IP 2017-02-10 ltm global-settings rule(1)

ltm global-settings traffic-control

NAME

traffic-control - Configures the global settings that pertain to traffic control for the BIG-IP(r) and VIPRION(r) local traffic management systems.

MODULE

ltm global-settings

SYNTAX

Configure the traffic-control component within the ltm global-settings module using the syntax shown in the following sections.

MODIFY

modify traffic-control

options:

accept-ip-options [disabled | enabled]
accept-ip-source-route [disabled | enabled]
allow-ip-source-route [disabled | enabled]
continue-matching [disabled | enabled]
max-icmp-rate [integer value: 0 ~ 2147483647]
max-reject-rate [integer value: 1 ~ 1000]
max-reject-rate-timeout [integer value: 0 ~ 300]
min-path-mtu [integer value: 68 ~ 1500]
path-mtu-discovery [disabled | enabled]
port-find-linear [integer value: 0 ~ 61439]
port-find-random [integer value: 0 ~ 1024]
port-find-threshold-warning [disabled | enabled]
port-find-threshold-trigger [integer value: 1 ~ 12]
port-find-threshold-timeout [integer value: 0 ~ 300]
reject-unmatched [disabled | enabled]

DISPLAY

list traffic-control

list traffic-control [option name]

show running-config traffic-control

show running-config traffic-control [option name]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the traffic-control component to modify how the system processes local traffic.

EXAMPLES

modify traffic-control accept-ip-options enabled
Specifies that the system accepts IPv4 packets with IP options.

list traffic-control
Displays the local traffic control global settings.

OPTIONS

accept-ip-options

Specifies whether the system accepts IPv4 packets with IP options. The default value is disabled.

accept-ip-source-route

Specifies whether the system accepts IPv4 packets with IP source route options that are destined for Traffic Management Microkernel (TMM). The default value is disabled.

To enable this option, you must also enable the accept-ip-options option.

allow-ip-source-route

Specifies whether the system allows IPv4 packets with IP source route options enabled to be routed through Traffic Management Microkernel (TMM). The default value is disabled.

To enable this option, you must also enable the accept-ip-options option.

continue-matching

Specifies whether the system matches against a less-specific virtual server when the more-specific one is disabled. When continue-matching is disabled, the default value, the system drops connections that request a disabled virtual server. In this case, the system rejects or drops packets depending on the value of the reject-unmatched option.

max-icmp-rate

Specifies the maximum rate per second at which the system issues Internet Control Message Protocol (ICMP) errors. The default value is 100 errors per second. The range is from 0 (zero) to 2147483647 errors per second. This option is useful for preventing ICMP-message storms.

max-reject-rate

Specifies the maximum rate per second at which the system issues reject packets (TCP RST or ICMP port unreachable). The default value is 250 per second. The range is from 1 to 1000 per second.

max-reject-rate-timeout

Specifies the time in seconds which the system ignores icmp port unreachable and tcp rst ratelimits on becoming active after a failover. The default value is 30 seconds. The range is from 0 to 300 seconds.

min-path-mtu

Specifies the minimum packet size that can traverse the path without suffering fragmentation, also known as path Maximum Transmission Unit(MTU). The default value is 296. The range is from 68 to 1500.

path-mtu-discovery

Specifies, when enabled, that the system discovers the maximum transmission unit (MTU) that it can send over a path, without fragmenting TCP packets. The default value is enabled.

port-find-linear

Specifies the maximum of ports to linearly search for outbound connections. The default value is 16. The range is from 0 to 61439.

port-find-random

Specifies the maximum of ports to randomly search for outbound connections. The default value is 16. The range is from 0 to 1024.

port-find-threshold-warning

Specifies if the ephemeral port-exhaustion threshold warning is to be monitored. The default is enabled.

port-find-threshold-trigger

Specifies the threshold warning's trigger which is the value of random port attempts when attempting to find an unused outbound port for a connection. The default is 8. The valid range is 1 - 12.

port-find-threshold-timeout

Specifies the threshold warning's timeout. This is the time in seconds since the last trigger value was hit and will drop the tuple if not hit. The default is 30 (1/2 minute) with range from 0 - 300.

reject-unmatched

Specifies, when enabled, that the system returns a TCP RESET or ICMP_UNREACH packet if no virtual servers on the system match the destination address of the incoming packet. When this option is disabled, the system silently drops the unmatched packet. The default value is enabled.

SEE ALSO

list, ltm node, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 ltm global-settings traffic-control(1)

Itm ifile

NAME

ifile - Configures an iFile.

MODULE

itm

SYNTAX

Configure the iFile component within the Itm module using the syntax shown in the following sections.

CREATE/MODIFY

create ifile [name]

modify ifile [name]

options:

app-service [[string] | none]

description [string]

file-name [ifile file object name]

edit ifile [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DISPLAY

list ifile

list ifile [[name] | [glob] | [regex]] ...]

show running-config ifile

show running-config ifile [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete ifile [name]

DESCRIPTION

You can use the ifile component to configure an iFile. The iFile can then be referenced from an iRule, to allow loading an external file into an iRule.

EXAMPLES

```
create ifile my_ifile file-name ifile_file_object_name
```

Creates an iFile named my_ifile, that gets its contents from the file object ifile_file_object_name.

```
list ifile all-properties
```

Displays all of the properties of all of the iFiles.

```
delete ifile my_ifile
```

Deletes the iFile named my_ifile.

OPTIONS

app-service

Specifies the name of the application service to which the iFile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the iFile. Only the application service can modify or delete the iFile.

description

User defined description.

file-name

The name of the iFile File Object that this iFile uses.

SEE ALSO

create, delete, edit, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 Itm ifile(1)

ltm lsn-log-profile

NAME

lsn-log-profile - Configures a Large-Scale Network Address Translation logging profile.

MODULE

ltm

SYNTAX

CREATE/MODIFY

```
create lsn-log-profile [name]
modify lsn-log-profile [name | all]
options:
  app-service [[string] | none]
  csv-format [disabled | enabled]
  start-outbound-session {
action [disabled | enabled | backup-allocation-only]
elements [add | delete | replace-all-with] {
  destination
}
}
  end-outbound-session {
action [disabled | enabled | backup-allocation-only]
elements [add | delete | replace-all-with] {
  destination
}
}
  start-inbound-session {
action [disabled | enabled | backup-allocation-only]
}
  end-inbound-session {
action [disabled | enabled | backup-allocation-only]
}
  quota-exceeded {
action [disabled | enabled ]
}
  errors {
action [disabled | enabled ]
}
  subscriber-id [disabled | enabled]
```

```
edit lsn-log-profile [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

```
reset-stats lsn-log-profile
```

```
reset-stats lsn-log-profile [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list lsn-log-profile
```

```
list lsn-log-profile [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config lsn-log-profile
```

```
show running-config lsn-log-profile [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
one-line
```

DELETE

```
delete lsn-log-profile [name | all]
```

DESCRIPTION

A LSN log profile allows fine grain control of the logging for LSN translation events. When attached to an LSN pool, you can control the events to enable logging, and the elements in the log entry.

EXAMPLES

```
create ltm lsn-log-profile my_lsn_log_profile end-inbound-session { action enabled } end-outbound-session {
action enabled }
```

Creates the LSN log profile `my_lsn_log_profile` that generates log entries for both inbound and session when the translation session ends.

```
delete lsn-log-profile my_lsn_log_profile
```

Deletes the LSN log profile named `my_lsn_log_profile`.

OPTIONS

`app-service`

Specifies the name of the application service to which this object belongs. The default value is `none`.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

`csv-format`

When enabled, use CSV log format for log entries. The default value is `disabled`.

events

The type of LSN translation events available for logging control.

start-outbound-session

Event for start of outbound translation session, when the outbound flow is created.

end-outbound-session

Event for end of outbound translation session, when the outbound flow is deleted.

start-inbound-session

Event for start of incoming connection to a translated address.

end-inbound-session

Event for end of incoming connection to a translated address.

quota-exceeded

Event for when client exceeded allocated resource limit.

errors

Event for when LSN encountered errors while attempting translation for clients.

subscriber-id

When enabled, the subscriber ID associated with a subscriber IP address will be printed in the logs.

action

Specify the logging action to be taken when a particular event is encountered.

enabled

Logging is enable for the event, whether translation is from the primary pool member or backup pool member.

backup-allocation-only

Logging is enable for the event, when translation is take from backup pool member only.

disabled

Logging is disable for the event.

elements

Optional elements that can be added to the log message.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

ltm lsn-pool, create, delete, edit, glob, list, ltm, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2017-02-15 ltm lsn-log-profile(1)

ltm lsn-pool

NAME

lsn-pool - Configures a Large-Scale Network Address Translation (or Carrier-Grade Network Address Translation) pool.

MODULE

ltm

SYNTAX

```
CREATE/MODIFY
create lsn-pool [name]
modify lsn-pool [name | all]
options:
  app-service [[string] | none]
  backup-members
    [add | delete | replace-all-with] {
[ip address/prefix length] ...
  }
  client-connection-limit [integer value]
  description [string]
```

```

egress-interfaces
  [add | delete | replace-all-with] {
[interface name] ...
  }
egress-interfaces-disabled
egress-interfaces-enabled
hairpin-mode [enabled | disabled]
icmp-echo [enabled | disabled]
inbound-connections [automatic | explicit | disabled]
log-publisher [log publisher name | none]
log-profile [log profile name | none]
members
  [add | delete | replace-all-with] {
[ip address/prefix length] ...
  }
mode [deterministic | napt | pba]
persistence {
  mode [none | address | address-port]
  timeout [integer]
}
pcp {
  profile [ name | none ]
  selfip [ name | none]
  dslite_tunnel [ name | none ]
}
port-block-allocation {
  block-idle-timeout [integer]
  block-lifetime [integer]
  block-size [integer]
  client-block-limit [integer]
  zombie-timeout [integer]
}
route-advertisement [enabled | disabled]
translation-port-range [integer low:integer high | integer]

```

```
edit lsn-pool [ [ [name] | [glob] | [regex] ] ... ]
```

```

options:
  all-properties
  non-default-properties

```

```
reset-stats lsn-pool
```

```
reset-stats lsn-pool [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list lsn-pool
```

```
list lsn-pool [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config lsn-pool
```

```
show running-config lsn-pool [ [ [name] | [glob] | [regex] ] ... ]
```

```

options:
  all-properties
  non-default-properties
  one-line

```

```
show lsn-pool
```

```
show lsn-pool [ [ [name] | [glob] | [regex] ] ... ]
```

```

options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  failure-cause

```

DELETE

```
delete lsn-pool [name | all]
```

DESCRIPTION

A large-scale NAT (LSN) pool is a set of networks and port numbers that the BIG-IP system uses as public-side addresses and ports. When you assign an LSN pool to a virtual server, the virtual server's clients have their private addresses (and/or ports) translated to a public address and/or port from the LSN pool. The public-side addresses and ports in the LSN pool are called translation addresses and ports.

EXAMPLES

```
create lsn-pool my_lsn_pool1 mode napt persistence { mode address-port timeout 600 } members add {
10.10.10.0/24 10.10.20.0/24 } translation-port-range 4000:5000 client-connection-limit 100
```

Creates the LSN pool `my_lsn_pool1` that contains the translation addresses in the range of (members) `10.10.10.0/24` and `10.10.20.0/24`, translation port range `4000-5000`, with a client connection limit of `100` connections per client. The translated address and port are persisted for `600` seconds. This LSN pool operates in NATP mode (Network Address and Port Translation mode), which is the default mode if not specified.

```
delete lsn-pool my_lsn_pool1
```

Deletes the LSN pool named `my_lsn_pool1`.

OPTIONS

```
app-service
```

Specifies the name of the application service to which this object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot

modify or delete this object. Only the application service can modify or delete this object.

backup-members

Specifies translation IP addresses available in the backup pool which is used by DNAT translation mode if DNAT mode translation fails and falls back to NAT mode. This is a collection of IP prefixes with their prefix lengths.

client-connection-limit

The maximum number of simultaneous translated connections a client or subscriber is allowed to have.

description

User defined description.

egress-interfaces

The set of interfaces on which the source address translation is allowed or disallowed. If egress-interfaces-enabled is specified, the source address translation is allowed only on the specified set of interfaces. If egress-interfaces-disabled is specified, source address translation is disabled on specified interfaces.

egress-interfaces-disabled

Source address translation is not allowed on the interfaces specified in the egress-interfaces set.

egress-interfaces-enabled

Source address translation is allowed on the interfaces specified in the egress-interfaces set.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

hairpin-mode

Enable or disable hairpinning for incoming connections.

When a client sends a packet to another client in the same private network, hairpin mode sends the packet directly to the destination client's private address; the BIG-IP system immediately translates the packet's public-side destination address. Rather than going out to the public network and coming back later for translation, the packet takes a hairpin turn at the BIG-IP device.

icmp-echo

Enable or disable ICMP echo on translated addresses.

inbound-connections

Modifies the inbound-connection mode for incoming connections to translation endpoints. A translation endpoint is the public-side address and port (X':x') for a private-side address (X:x). You can allow one of three algorithms for managing inbound connections:

Automatic

creates inbound mappings automatically from outbound traffic and allows inbound connections. Consider an outbound mapping from X:x to X':x'. If a connection comes from X:x through X':x', the BIG-IP system automatically creates a reverse mapping from X':x' back to X:x. A public-side station can respond through the X':x' address. This allows the BIG-IP system to provide Endpoint Independent Filtering (EIF) as defined in section 5 of RFC 4787
().

Explicit

only allows inbound connections for mappings that are explicitly created by another party, such as iRules or a PCP request. For example, if a PCP request creates a mapping of X:x to X':x' and the client at X:x uses it, an external caller can respond to the client through X':x'. However, if a client at M:m automatically makes a NAT'ed connection through M':m', the BIG-IP does not support an inbound connection from M':m' back to M:m.

Disabled

disables inbound connections to translation end-points (X':x'). If there is a mapping of X (a private-side IP address) to X' (a public-side IP), connections can only go out from X through X'. If a public-side recipient tries to answer at the client's public-side X' address, the BIG-IP system does not map X' back to X. The inbound connection never happens.

Port Control Protocol (PCP) is not supported if you use this setting.

log-publisher

Specify the name of the log publisher which logs translation events. See help sys log-config for more details on the logging sub-system. Use the "sys log-config publisher" component to set up a log publisher.

log-profile

Specify the name of the LSN log profile which controls the logging of translation events. See help ltm lsn-log-profile for more details on the logging profile sub-system. Use the "ltm lsn-log-profile profile" component to set up a LSN log profile.

members

Specifies the set of translation IP addresses available in the pool. This is a collection of IP prefixes with their prefix lengths. All public-side addresses come from the subnets you enter in this property.

mode Specifies which kind of translation address mapping is performed when an address is translated. Available options are NAT, Deterministic, and PBA.

NAPT (Network Address Port Translation) assigns translation addresses and ports in round-robin fashion. The algorithm first cycles through translation addresses and then through translation ports.

Deterministic (DNAT) is a reversible translation method. A given client address and port always translates to a particular public address and port from the LSN pool. This method has the following restrictions:

it is only available for NAT44 translations,
it does not support connections through DS-Lite tunnels,
subscriber connections must be received over a VLAN with the property, `cmp-hash`, set to "source ip,"
the egress to the Internet must be over a VLAN with the property, `cmp-hash`, set to "dest ip,"
any virtual server ("itm virtual") that uses this LSN pool must have a source property set to an IP prefix containing fewer than 231 addresses. For example, the source cannot be 0.0.0.0/0.

PBA (Port Block Allocation) assigns 'blocks' of the translation addresses and ports to individual clients. All client connections are restricted to the allocated port blocks. Only block allocations and deallocations are logged in order to reduce the volume of logs.

subscriber connections must be received over a VLAN with the property, `cmp-hash`, set to "source ip,"
the egress to the Internet must be over a VLAN with the property, `cmp-hash`, set to "dest ip,"

You can access your VLAN configurations through the "net vlan" component. You can find the VLANs used by your virtual server by showing or listing the "itm virtual" component.

`name` Specifies a unique name for the `lsn-pool` component. This option is required for the commands `create`, `delete`, and `modify`.

`persistence`
Configure the persistence settings for LSN translation entries. Persistence is the preservation of a public-side IP address for a client from session to session.

`persistence.mode`
Configure the persistence mode for LSN translation entries. You can enter `address`, `address-port`, or `none`.

`address`
causes the BIG IP software to attempt to keep the IP address persistent but not necessarily the port. If a client's private IP address:port combination is X:x, it's public-side address may be X':a in one session, X':b in the next session, X':c in a third session, and so on.

`address-port`
causes the BIG IP software to attempt to keep the IP address and port persistent. If a client's private IP address:port combination is X:x, and it's public-side address is X':x' in the first session, it remains X':x' in all future sessions.

This is called "Endpoint Independent Mapping" in RFC 4787 ().

This is the only supported setting for PCP, which you configure with the `pcp` property.

`none` prevents the BIG IP software from attempting any IP address or port persistence. An address:port combination of X:x is never guaranteed to have the same public-side address or port in two sessions.

`persistence.timeout`
After the most-recent session where address:port X:x translated to X':x' on the public side, a timer begins. If the timer expires before X:x has another session, X' or x' may be used as the public side of another address:port. Use this parameter to set the timeout (in seconds) for address and port persistence.

`pcp` A Port Control Protocol (PCP) client can set (or at least learn) its own translation (public-side) IP address and/or port. It can also set the address and/or port of a third-party client. PCP is defined in RFC 6887 (see).

`pcp.profile`
Specifies the PCP profile to use for this LSN pool. This PCP profile defines the settings to use for communication with PCP clients. Use the `create itm profile pcp` command to create a new PCP profile.

PCP requires a profile (defined with this property) and either a `pcp.selfip` or `pcp.dslite tunnel` where clients can send their PCP requests.

If you remove this profile option, you must specifically remove any `pcp.selfip` or `pcp.dslite tunnel`, too.

`pcp.selfip`
Specifies the PCP Server self-IP address for this LSN pool. The virtual server's clients send their PCP packets to this address. Use the `create net self` command to create a self-IP address, then use that address for this parameter. Choose a self-IP address in a VLAN that is reachable by the virtual server's clients.

`pcp.dslite`
Specifies a DS-LITE tunnel for PCP packets. Whenever a client sends a PCP packet through this tunnel, the BIG-IP device uses the PCP profile you choose with the `pcp.profile` property.

A DS-LITE tunnel places each IPv4 packet into the payload of an IPv6 packet. The IPv6 packet carries the IPv4 packet between customer equipment and the BIG-IP system, which then removes the IPv4 packet, uses NAT to translate its IPv4 addresses, and sends it to its destination.

You cannot use this property if the `mode` property is set to `Deterministic`.

`port-block-allocation`
Configures the port block settings for PBA mode.

`port-block-allocation.block-idle-timeout`
Configures the time after the last connection using the block is freed that the block assignment expires.

The default value is 3600 seconds.

`port-block-allocation.block-lifetime`

Configures the timeout after which the block is no longer used for new port allocations. The block becomes a zombie block. The default is 0 which corresponds to an infinite timeout.

`port-block-allocation.block-size`

Configures the number of ports in a block. The default value is 64.

`port-block-allocation.client-block-limit`

Configures the number of blocks that can be assigned to a single subscriber IP address. The default value is 1.

`port-block-allocation.zombie-timeout`

Configures the timeout after which connections using the zombie block are killed. After connections are killed zombie block is freed after `port-block-allocation.block-idle-timeout`. This parameter is unused unless the `port-block-allocation.block-lifetime` is set. The default value is 0 which corresponds to infinite timeout.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`route-advertisement`

Specifies whether route advertisement is enabled or disabled for translated IP addresses.

`translation-port-range`

Specifies the range of port numbers available for use with translation IP addresses.

`failure-cause`

Displays the failure-cause table for this lsn-pool. There are many different possible failure causes and only the failures that occur will be displayed. This information can be useful for determining why a translation is failing.

SEE ALSO

ltn profile `pcp`, ltn virtual, net self, net vlan, create, delete, edit, glob, list, ltn, modify, regex, reset-stats, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-07-27 ltn lsn-pool(1)

ltn message-routing diameter peer

NAME

`peer` - Configures a peer for routing Diameter protocol messages.

MODULE

ltn message-routing diameter

SYNTAX

Configure the peer component within the ltn message-routing diameter module using the syntax shown in the following sections.

CREATE/MODIFY

`create peer [name]`

`modify peer [name]`

options:

`app-service` [[string] | none]

`auto-initialization` [enabled | disabled]

`auto-initialization-interval` [integer]

`connection-mode` [per-peer | per-blade | per-tmm | per-client | per-client-per-blade |

per-client-per-tmm | per-peer-alternate-tmm | per-client-alternate-tmm]

`description` [string]

`number-connections` [integer]

`pool` [name]

`ratio` [integer]

`transport-config` [transport-config]

`edit peer [[name] | [glob] | [regex]] ...]`

options:

`all-properties`

DISPLAY
list peer
list peer [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 one-line

DELETE
delete peer [name]

DESCRIPTION

You can use the peer component to manage a named Diameter peer. A peer specifies the pool for the Diameter router to use as the destination for Diameter routes. You can also use the peer component to specify how many connections the parser creates to a remote host and what transport the parser uses to establish the connection.

EXAMPLES

```
create peer my_peer { pool my_pool transport { type virtual name my_vip }
```

Creates a Diameter peer named `my_peer` which uses the settings of `my_vip` to establish a connection with a pool member from pool `my_pool`.

OPTIONS

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`auto-initialization`

If enabled, the BIGIP will automatically create outbound connections to the active pool members in the specified pool using the configuration of the specified transport-config. For auto-initialization to attempt to create a connection, the peer must be included in a route that is attached to a router instance. For each router instance that the peer is contained in, a connection will be initiated. The auto-initialization logic will verify at a configurable interval if the a connection exists between the BIG-IP and the pool members of the pool. If a connection does not exist, it will attempt to reestablish one. The default is disabled.

`auto-initialization-interval`

Specifies the interval (in milliseconds) that attempts to initiate a connection occur. Valid ranges are from 500ms to 65535ms. The default is 5000ms.

`connection-mode`

Specifies how the number of connections per host is limited. Note a host (specified in the referred pool) may exist more than one peer object, and those peer objects may have different settings for `connection-mode` and `number_connections`. Thus, these settings specify how messages routed through this peer are distributed between a set of connections, not the maximum number of connections to a specified host. The default value is `per-peer`.

`per-blade`

Specifies the number of connections to a remote host per blade in the cluster.

`per-client`

Specifies the number of connections to a remote host per client connection.

`per-client-per-blede`

Specifies the number of connections to a remote host per client connection for each blade.

`per-client-per-tmm`

Specifies the number of connections to a remote host per client connection for each tmm.

`per-peer`

Specifies the number of connections to a remote host.

`per-peer-alternate-tmm`

Specifies the number of connections to a remote host. Any new connections will be opened on an alternate TMM to spread CPU usage.

`per-client-alternate-tmm`

Specifies the number of connections to a remote host per client connection. Any new connections will be opened on an alternate TMM to spread CPU usage.

`per-tmm`

Specifies the number of connections to a remote host per TMM in the system.

`description`

User defined description.

`number-connections`

Specifies the distribution of connections between the BIG-IP system and a remote host. The default value is 1.

`pool` Specifies the name of the pool to which the Diameter parser routes messages.

`ratio`

Specifies the ratio the Diameter router uses to select a peer from a list of peers for the ltm message-routing diameter route. The default value is 1.

transport-config

Specifies the name of the transport configuration (ltm message-routing diameter transport-config) the message router uses to create an outgoing connection.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh, ltm message-routing diameter route ltm message-routing diameter profile session

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2019-11-20 ltm message-routing diameter peer(1)

ltm message-routing diameter profile router

NAME

router - Configures a Diameter Router profile.

MODULE

ltm message-routing diameter profile

SYNTAX

Configure the router component within the ltm message-routing diameter profile module using the syntax shown in the following sections.

CREATE/MODIFY

create router [name]

modify router [name]

options:

app-service [[string] | none]

associate-clientside-to-poolmember [disabled | enabled]

defaults-from [[name] | none]

description [string]

ha-message-sweeper-interval [integer]

ignore-peer-port [disabled | enabled]

irule-scope-message [yes | no]

max-pending-bytes [integer]

max-pending-messages [integer]

max-retries [integer]

mirrored [disabled | enabled]

pending-request-sweeper-interval [integer]

per-peer-stats [enabled | disabled]

routes [add | default | delete | none | replace-all-with] {
[route_name] ...
}

supported-applications [[integer] ...]

traffic-group [[string] | default | non-default | none]

transaction-timeout [integer]

use-local-connection [disabled | enabled]

edit router [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats router

reset-stats router [[[name] | [glob] | [regex]] ...]

DISPLAY

list router

list router [[[name] | [glob] | [regex]] ...]

show running-config router

show running-config router [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show router

show router [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete router [name]

DESCRIPTION

You can use the router component to manage a Diameter router profile.

EXAMPLES

```
create router my_router_profile defaults-from router
```

Creates a Diameter router profile named my_router_profile using the system defaults.

```
create router my_router_profile routes add { route1 route2 }
```

Creates a Diameter profile named my_router_profile with two static routes.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

associate-clientside-to-poolmember

If enabled, the configured routes will be scanned for a pool member that matches a new clientside connection. If found, the clientside connection will be associated with the pool member allowing the activity on the clientside connection to be included with the activity of the pool member. The default value is disabled.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is router.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ha-message-sweeper-interval

Specifies the maximum time in milliseconds that a message will be held on the standby device as it waits for the active device to route the message. Messages on the standby device held for longer than the configurable sweeper interval, will be dropped. The default value is 1000 milliseconds.

ignore-peer-port

If selected (enabled), any connection from a configured peer will be suitable for routing a message to, regardless of its remote port number. The default value is enabled.

irule-scope-message

If set to yes, iRule events are scoped to the message executing (each message has its own execution context). This allows multiple messages to process iRules concurrently without waiting for commands from other messages. If set to no, all Diameter iRule events are scoped to the connection flow (that is all share a single execution context per flow) and therefore are forced to execute one at a time even though they are processing independent messages (this is the legacy mode). The default value is no. Some existing iRules might need to be adjusted to use the per-message scope.

max-pending-bytes

Limits the number of bytes contained within messages held pending while waiting for a connection to a peer. If irule-scope-message is set to yes, also limits the number of bytes contained within messages that may concurrently process iRule events. Once reached any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 65536.

max-pending-messages

Limits the number of messages held pending while waiting for a connection to a peer. If irule-scope-message is set to yes, also limits the number of messages that may concurrently process iRule events. Once the limit is reached, any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 64.

max-retries

This attribute sets the maximum number of time a message may be resubmitted for rerouting by the MR::retry iRule command. The default value is 1.

mirrored

If enabled, connection created on the active device of the traffic-group specified, will be mirrored on the standby device. Messages processed on the active device will also be mirrored and perform equivalent processing on the standby device.

partition

Displays the administrative partition within which the component resides.

per-peer-stats

If enabled, the profile specific statistics will be captured for each pool member. The default value is disabled.

pending-request-sweeper-interval

Specifies the interval in milliseconds between passes of the pending request sweeper. The pending request sweeper will delete pending request entries that are older than twice the transaction-timeout. If set to 0 the sweeper will be disabled. The default value is 60000ms.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

routes

Specifies a list of static routes. The ordering of the route entries is insignificant.

supported-applications

Defines the list of application ID values that will be supported. If a message is received with an Auth-Application-ID or Acct-Application-ID AVP that does not match any of the values in this list, it will be dropped and the connection will be reset. A value of 0 will match all application IDs. The default value is none.

traffic-group

Specifies the traffic group on which the router is active. The default traffic group is inherited from the containing folder.

inherited-traffic-group

Read-only property that indicates if the traffic-group is inherited from the parent folder.

transaction-timeout

Specifies the maximum time (in seconds) between a request and its response. A provisional response restarts the timer. The default value is 10 seconds. Note: This may not affect all transactions. The scenarios where the system waits for response (eg. a final response for REGISTER request), are impacted by dropping any persistent data maintained for the request.

use-local-connection

Enables or disables a preference for local connections established by the ingress TMM over connections established by other TMM's when selecting the egress connection to destination peer. By default this attribute is enabled.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016. All rights reserved.

BIG-IP 2018-07-10 ltm message-routing diameter profile router(1)

ltm message-routing diameter profile session

NAME

session - Configures a Diameter Session profile.

MODULE

ltm message-routing diameter profile

SYNTAX

Configure the session component within the ltm message-routing diameter profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create session [name]
modify session [name]
options:
  acct-application-id [integer]
  app-service [[string] | none]
  array-acct-application-id [[list of integers] | none]
  array-auth-application-id [[list of integers] | none]
  array-retransmission-result-codes [[list of integers] | none]
  auth-application-id [integer]
  defaults-from [[name] | none]
  description [string]
  dest-host-rewrite [string]
  dest-realm-rewrite [string]
  disconnect-peer-action [disable | force-offline | none]
  dynamic-route-insertion [disabled | enabled]
  dynamic-route-lookup [disabled | enabled]
  dynamic-route-timeout [integer]
  discard-unroutable [disabled | enabled]
  egress-critical-message-rate-limit [integer]
  egress-major-message-rate-limit [integer]
  handshake-timeout [integer]
  host-ip-address [disabled | enabled]
  ingress-critical-message-rate-limit [integer]
```

ingress-major-message-rate-limit [integer]
loop-detection [disabled | enabled]
max-message-size [integer]
max-retransmissions [integer]
max-watchdog-failures [integer]
origin-host [string]
origin-host-rewrite [string]
origin-realm [string]
origin-realm-rewrite [string]
egress-critical-message-rate-limit [integer]
persist-avp [string]
persist-timeout [integer]
persist-type [avp | custom | none]
product-name [string]
reset-on-timeout [disabled | enabled]
respond-unroutable [disabled | enabled]
retransmission-action [disabled | busy | unable | retransmit | retransmit-alternate]
retransmission-queue-limit-low [integer]
retransmission-queue-limit-high [integer]
retransmission-queue-max-bytes [integer]
retransmission-queue-max-messages [integer]
retransmission-timeout [integer]
route-unconfigured-peers [disabled | enabled]
vendor-id [integer]
vendor-specific-vendor-id [integer]
vendor-specific-acct-application-id [integer]
vendor-specific-auth-application-id [integer]
watchdog-timeout [integer]

edit session [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

reset-stats session

reset-stats session [[[name] | [glob] | [regex]] ...]

DISPLAY

list session

list session [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line

show running-config session

show running-config session [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line

show session

show session [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE

delete session [name]

DESCRIPTION

You can use the session component to manage a Diameter session profile.

EXAMPLES

```
create session my_session_profile defaults-from session
```

Creates a Diameter session profile named my_session_profile using the system defaults.

```
create session my_session_profile { reset-on-timeout disabled }
```

Creates a Diameter profile named my_session_profile that will not reset the connection when watchdog failure exceed maximum-watchdog-failures.

OPTIONS

acct-application-id

Specifies as an integer the Accounting identifier for specific application, as specified in RFC 6733.

This value will be appended at the end of array-acct-application-id in capabilities exchange messages if it doesn't already exist in it.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

array-acct-application-id

Specifies as a whitespace separated list of integers the Accounting identifier(s) for specific

application(s), as specified in RFC 6733.

array-auth-application-id

Specifies as a whitespace separated list of integers the Authentication and Authorization identifier(s) for specific application(s), as specified in RFC 6733.

array-retransmission-result-codes

Specifies as a whitespace separated list of integers that define result codes that if received in an answer message will trigger retransmission.

auth-application-id

Specifies as an integer the Authentication and Authorization identifier for specific application, as specified in RFC 6733. This value will be appended at the end of array-auth-application-id in capabilities exchange messages if it doesn't already exist in it.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is diametersession.

description

User defined description.

dest-host-rewrite

Specifies the destination host AVP to which the specified value on the egress will be rewritten.

dest-realm-rewrite

Specifies the destination realm AVP to which the specified value on the egress will be rewritten.

discard-unroutable

When selected (enabled), messages that do not match any known route will be silently discarded. When disabled, unroutable messages are routed back to the connection where they came from. The default value is enabled.

disconnect-peer-action

Specifies the state of peer based on Disconnect Peer Request received from peer. The default value is none. The options are:

none Terminates connection on receiving DPR. Connection can be re-established between peer and BIGIP.

disable

A node continues to process persistent and active connections. It can accept new connections only if the connections belong to an existing persistent session.

force-offline

A node allows existing connections to time out, but no new connections are allowed.

egress-critical-message-rate-limit

If the number of messages egressed to a peer (pool member) in a second exceeds the provided limit, a SNMP trap, The number of messages sent to a peer is above the critical rate limit threshold, will be sent. If the number of messages egressed to a peer (pool member) in a second drops below the provided limit, a SNMP trap, The number of messages sent to a peer is back under the critical rate limit threshold, will be sent. A value of 0 will disable the SNMP trap. The default value is 0.

egress-major-message-rate-limit

If the number of messages egressed to a peer (pool member) in a second exceeds the provided limit, a SNMP trap, The number of messages sent to a peer is above the major rate limit threshold, will be sent. If the number of messages egressed to a peer (pool member) in a second drops below the provided limit, a SNMP trap, The number of messages sent to a peer is back under the major rate limit threshold, will be sent. A value of 0 will disable the SNMP trap. The default value is 0.

dynamic-route-insertion

Specifies whether dynamic route insertion is enabled for this Diameter session profile. If enabled, routes will be added to route incoming messages toward the connected peer, by its origin-host. The default value is disabled.

dynamic-route-lookup

Specifies whether dynamic route lookup is enabled for this Diameter session profile. If enabled, the destination-host of messages received via this profile will be used to find a route added from connections with dynamic-route-insertion enabled. The default value is disabled.

dynamic-route-timeout

Specifies how long after a connection is closed will the dynamic route be deleted from the route table. The default value is 300

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

handshake-timeout

Specifies the number of seconds before the peer handshake times out. The default is 10 seconds.

host-ip-address

Specifies the value that will be used in the Host-IP-Address AVP sent in Capabilities-Exchange-Request and Capabilities-Exchange-Answer messages. When unset (default), the Diameter router will use the virtual server's IP address.

ingress-critical-message-rate-limit

If the number of messages received from a peer (pool member) in a second exceeds the provided limit, a

SNMP trap, The number of messages from a peer is above the critical rate limit threshold, will be sent. If the number of messages received from a peer (pool member) in a second drops below the provided limit, a SNMP trap, The number of messages from a peer is back under the critical rate limit threshold, will be sent. A value of 0 will disable the SNMP trap. The default value is 0.

ingress-major-message-rate-limit

If the number of messages received from a peer (pool member) in a second exceeds the provided limit, a SNMP trap, The number of messages from a peer is above the major rate limit threshold, will be sent. If the number of messages received from a peer (pool member) in a second drops below the provided limit, a SNMP trap, The number of messages from a peer is back under the major rate limit threshold, will be sent. A value of 0 will disable the SNMP trap. The default value is 0.

loop-detection

Specifies whether loop detection will be performed on requests received by this session profile. The default value is enabled. When set, the Diameter session profile will reject messages that it has already seen. See RFC 6733 section 6.1.3.

max-message-size

Specifies the maximum number of bytes acceptable in a Diameter message. The default value is 0 which indicates that there is no message size restriction for this session. Note: Message size is also restricted by the database variable "diameter.message.maxlen"; the smallest value is used as a maximum. Messages exceeding this size are silently discarded.

max-retransmissions

Specifies the maximum number of retransmissions of a Diameter message. The default value is 0.

max-watchdog-failures

Specifies the maximum number of device watchdog failures that the traffic management system can receive before it tears down the connection. After the system receives this number of device watchdog failures, it closes the connection. The default value is 1.

origin-host

Specifies the identifier of the originating server in the form siteserver.f5.com. Must specify the origin-host.

origin-host-rewrite

Specifies the value to rewrite to the Origin-Host AVP on egress.

origin-realm

Specifies the Origin-Realm AVP data. Must specify the origin-realm.

origin-realm-rewrite

Specifies the value to rewrite to the Origin-Realm AVP on egress.

peer-delay-critical-limit

If the average peer delay exceeds the provided limit, a SNMP trap, PeerHealth exceeds critical, will be generated. If the average peer delay drops below the provided limit, a SNMP trap, PeerHealth back under critical, will be generated. A value of 0 will disable the SNMP trap. The default value is 0.

persist-avp

Specifies the Diameter AVP that is used for persistence. The format is avp[index] for a single AVP or a[x]:b[y]:c[z]:d[w] for a grouped AVP. There may be at most 4 AVPs in a group. The AVP name is used as the session-key; it may be an ASCII string or numeric ID in the range 1 to 4294967295 (AVP code can be specified instead of AVP name). Note: The default value is "SESSION-ID[0]". A grouped-avp can be specified with the following syntax: grouped-avp-name[index]:nested-avp1[index1]:nested-avp2[index2], where "nested-avp1" and "nested-avp2" are the AVPs in the grouped AVP.

persist-timeout

Specifies the timeout value (in seconds) for persistence entries. The default value is 180. Note: Its recommended to have the persist-timeout to be greater than transaction timeout, specified in the Diameter router configuration, as the lesser of the two is used when creating the persist record on receiving of the first Diameter request message. Upon receiving of the response for the first Diameter request message the persistence record is updated with the persist-timeout value. For any subsequent responses received the persist timeout is updated for the persist record.

persist-type Specifies the type of the persistence. The options are:

avp Persist based on avp in the message.

custom

Persist based on a custom key set using iRule.

none Persistence is disabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reset-on-timeout

When enabled, the system resets the connection when the number of watchdog failures exceeds the value of max-watchdog-failures. The default value is enabled.

respond-unroutable

When selected (enabled), messages that do not match any known route will be transformed into an error answer message and sent to the originator of the request. When disabled, unroutable request messages are routed back to the connection where they came from. The default value is disabled.

retransmission-action

Specifies the action performed when retransmission has been triggered for a request message. The options are:

disabled

Retransmission is disabled. This is the default action.

busy An answer message is generated with a TOO_BUSY result code and returned to the originator of the request.

unable

An answer message is generated with a UNABLE_TO_DELIVER result code and returned to the originator of the request.

retransmit

The request message will be retransmitted.

retransmit-alternate

The request message can be retransmitted to a different pool member.

retransmission-queue-limit-high

Specifies the high watermark for the retransmission queue (in percentage). If the retransmission queue exceeds this limit the transport window will begin closing. A value of 0 will disable closing the transport window. Valid range from 0 to 100. The default value is 90.

retransmission-queue-limit-low

Specifies the low watermark for the retransmission queue (in percentage). If the retransmission queue drops below this limit the transport window will reopen. Valid range from 0 to 100. The default value is 60.

retransmission-queue-max-bytes

Specifies the maximum number of bytes that can be stored in a connections retransmission queue. A value of 0 will disable this limit. The default value is 131072 bytes.

retransmission-queue-max-messages

Specifies the maximum number of messages that can be stored in a connections retransmission queue. A value of 0 will disable this limit. The default value is 1024 messages.

retransmission-timeout

Specifies the timeout for retransmission of a Diameter request (in seconds). A value of 0 will disable the retransmission timer. The default value is 10 seconds.

route-unconfigured-peers

When enabled, all connections will be allowed. When disabled, connections from peers whose IP addresses cannot be found in a statically configured route will be rejected. The default value is enabled.

vendor-id

Specifies the vendor identification number assigned to your diameter server by the Internet Assigned Numbers Authority (IANA). The default value is 3375.

vendor-specific-vendor-id

Specifies the vendor ID number that will be sent in Vendor-Specific-Application-ID AVPs. A value of 0 disables the feature. If this value is set, exactly one of either vendor-specific-acct-app-id or vendor-specific-auth-app-id must also be specified. The default value is 0.

vendor-specific-acct-app-id

Specifies the accounting application ID number that will be sent in Vendor-Specific-Application-ID AVPs. A value of 0 disables the feature. If this value is set, vendor-specific-vendor-id must be set and vendor-specific-auth-app-id must be unset. The default value is 0.

vendor-specific-auth-app-id

Specifies the authentication/authorization application ID number that will be sent in Vendor-Specific-Application-ID AVPs. A value of 0 disables the feature. If this value is set, vendor-specific-vendor-id must be set and vendor-specific-acct-app-id must be unset. The default value is 0.

watchdog-timeout

Specifies the watchdog timeout in seconds. This setting specifies the number of seconds that a connection is idle before the device watchdog request is sent. A value of 0 means BIG-IP will not send a device watchdog request to either client or server side. The default value is 10 seconds.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2015. All rights reserved.

ltm message-routing diameter route

NAME

route - Configures a static route for use in Diameter protocol message routing.

MODULE

ltm message-routing diameter

SYNTAX

Configure the route component within the ltm message-routing diameter module using the syntax shown in the following sections.

CREATE/MODIFY

create route [name]

modify route [name]

options:

app-service [[string] | none]

application-id [integer]

description [string]

destination-realm [[string] | none]

origin-realm [[string] | none]

peer-selection-mode [ratio | sequential]

peers { [none] | [peer_name ...] }

virtual-server [virtual-server_name]

edit route [[[name] | [glob] | [regex]] ...]

options:

all-properties

DISPLAY

list route

list route [[[name] | [glob] | [regex]] ...]

show running-config route

show running-config route [[[name] | [glob] | [regex]] ...]

options:

all-properties

one-line

DELETE

delete route [name]

DESCRIPTION

You can use the route component to define origin and destination realms, virtual server, peers, and peer selection mode of a message routing Diameter static route.

EXAMPLES

```
create route my_route
```

Creates a route instance named my_route using the system defaults.

```
create route my_route peers { peer1 peer2 }
```

Creates a route instance named my_route that will use two peers for forwarding messages.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

application-id

Specifies the application-id to match in the Diameter message. Default value 0 matches every application-id.

description

User defined description.

destination-realm

When specified, match the Destination-Realm AVP value in the message. Default value of "" specifies all destination-realms may be routed.

origin-realm

When specified, match the Origin-Realm AVP value in the message. Default value of "" specifies all origin-realms may be routed.

peer-selection-mode

Specifies the mode of selecting a peer from a list of peers. The options are:

ratio

Peers are selected based on their weights in comparison with other peers.

sequential

Peers are selected in the order listed. All traffic will route the first peer unless all pool

members in the peer are marked down.

peers

Specifies an ordered list of peers to use for forwarding messages.

virtual-server

Restricts routing for this route to connections originating on the specified virtual server. The default value is none which means the route is not restricted and messages originating on any connection may be routed to the route.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, show, tmsh, ltm message-routing diameter route

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-03-14 ltm message-routing diameter route(1)

ltm message-routing diameter transport-config

NAME

transport-config - Configures a Diameter transport-config instance for routing Diameter message protocol messages.

MODULE

ltm message-routing diameter

SYNTAX

Configure the transport-config component within the ltm message-routing diameter module using the syntax shown in the following sections.

CREATE/MODIFY

```
create transport-config [name]
modify transport-config [name]
options:
  app-service [[string] | none]
  description [string]
  profiles [add | delete | replace-all-with] {
    [profile_name ...] {
context [all | clientside | serverside]  read-only attribute for v12.0.0 or greater.
    }
  }
  rules { [none | [rule_name ... ] ] }
  source-address-translation {
  options:
  pool [ [pool_name] | none]
  type [ automap | snat | none ]
  }
  source-port [integer]
  source-port-mode [change | preserve | preserve-strict]
```

```
edit transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list transport-config
list transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete transport-config [name]
```

DESCRIPTION

You can use the transport-config component to define the profiles, rules, and source-address-translation of an outgoing connection.

EXAMPLES

```
create transport-config my_transport-config
```

Creates a transport-config instance named `my_transport-config` using the system defaults.

```
create transport-config my_transport-config { profiles add { my_diameter my_tcp } }
```

Creates a transport-config instance named `my_transport-config` that will use two profiles, `my_diameter` and `my_tcp`, to create and configure an outgoing connection. The outgoing connection is automatically configured with the router instance that created the connection.

OPTIONS

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`description`

User defined description.

`profiles`

Specifies a list of profiles that the outgoing connection uses to use to direct and manage traffic. The default value is none.

`rules`

Specifies a list of iRules, separated by spaces, that customize the transport configuration to direct and manage traffic. The default value is none.

`source-address-translation`

Specifies the type of source address translation enabled for the transport configuration, as well as the pool that the source address translation uses.

`pool` Specifies the name of a SNAT pool used by the specified transport configuration.

`type` Specifies the type of source address translation associated with the specified transport configuration.

The options are:

`automap`

Specifies the use of self IP addresses for transport configuration server source address translation.

`none` Specifies no source address translation is used by the transport configuration.

`snat` Specifies the use of a SNAT pool of translation addresses for virtual server source address translation.

`source-port`

Specifies the source port to be used for the connection being created. If `source-port-mode` is `change`, this setting has no effect. The default value is 0.

`source-port-mode`

Specifies how the system should select a source port for the outgoing connection. The default value is `change`.

The options are:

`change`

Selects an ephemeral source port for the outgoing connection.

`preserve`

Attempts to use the value of `source-port` for the outgoing connection, if specified. Otherwise attempts to preserve the source port of the incoming connection.

`preserve-strict`

Forces the outgoing connection to use the value of `source-port`, if specified. Otherwise forces the new connection to preserve the source port of the incoming connection. The system will fail to create a new outgoing connection if the specified source port is already in use.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `regex`, `tmsl`, `ltm message-routing diameter route`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2019-10-07 ltm message-routing diameter transport-config(1)

NAME

peer - Configures a peer for routing generic message protocol messages.

MODULE

ltm message-routing generic

SYNTAX

Configure the peer component within the ltm message-routing generic module using the syntax shown in the following sections.

CREATE/MODIFY

create peer [name]

modify peer [name]

options:

app-service [[string] | none]

auto-initialization [enabled | disabled]

auto-initialization-interval [integer]

connection-mode [per-peer | per-blade | per-tmm | per-client | per-client-per-blade |

per-client-per-tmm | per-peer-alternate-tmm | per-client-alternate-tmm]

description [string]

number-connections [integer]

pool [name]

ratio [integer]

transport-config [transport-config]

edit peer [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list peer

list peer [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete peer [name]

DESCRIPTION

You can use the peer component to manage a named generic message peer. A peer specifies the pool for the generic message parser to use as the destination for generic message routes. You can also use the peer component to specify how many connections the parser creates to a remote host and what transport the parser uses to establish the connection.

EXAMPLES

```
create peer my_peer { pool my_pool transport { type virtual name my_vip } }
```

Creates a generic message peer named my_peer which uses the settings of my_vip to establish a connection with a pool member from pool my_pool.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auto-initialization

If enabled, the BIGIP will automatically create outbound connections to the active pool members in the specified pool using the configuration of the specified transport-config. For auto-initialization to attempt to create a connection, the peer must be included in a route that is attached to a router instance. For each router instance that the peer is contained in, a connection will be initiated. The auto-initialization logic will verify at a configurable interval if the a connection exists between the BIG-IP and the pool members of the pool. If a connection does not exist, it will attempt to reestablish one. The default is disabled.

auto-initialization-interval

Specifies the interval (in milliseconds) that attempts to initiate a connection occur. Valid ranges are from 500ms to 65535ms. The default is 5000ms.

connection-mode

Specifies how the number of connections per host is limited. Note a host (specified in the referred pool) may exist more than one peer object, and those peer objects may have different settings for connection-mode and number_connections. Thus, these settings specify how messages are routed through this peer are distributed between a set of connections, not the maximum number of connections to a specified host. The default value is per-peer.

per-peer

Specifies the number of connections to a remote host.

per-blade

Specifies the number of connections to a remote host per blade in the cluster.

per-tmm

Specifies the number of connections to a remote host per TMM in the system.

per-client

Specifies the number of connections to a remote host per client connection.

per-client-per-blede

Specifies the number of connections to a remote host per client connection for each blade.

per-client-per-tmm

Specifies the number of connections to a remote host per client connection for each tmm.

per-peer-alternate-tmm

Specifies the number of connections to a remote host. Any new connections will be opened on an alternate TMM to spread CPU usage.

per-client-alternate-tmm

Specifies the number of connections to a remote host per client connection. Any new connections will be opened on an alternate TMM to spread CPU usage.

description

User defined description.

number-connections

Specifies the distribution of connections between the BIG-IP system and a remote host. The default value is 1.

pool Specifies the name of the pool to which the generic parser routes messages.

ratio

Specifies the ratio the generic message parser uses to select a peer from a list of peers for the ltm message-routing generic route. The default value is 1.

transport-config

Specifies the name of the transport configuration (ltm message-routing generic transport-config) the message router uses to create an outgoing connection.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh, ltm message-routing generic route ltm message-routing generic protocol

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2014. All rights reserved.

BIG-IP 2019-11-20 ltm message-routing generic peer(1)

ltm message-routing generic protocol

NAME

protocol - Configures a generic message protocol component for parsing generic messages.

MODULE

ltm message-routing generic

SYNTAX

Configure the protocol component within the ltm message-routing generic module using the syntax shown in the following sections.

CREATE/MODIFY

create protocol [name]

modify protocol [name]

options:

defaults-from [[name] | none]

description [string]

disable-parser [yes | no]

max-egress-buffer [integer]

max-message-size [integer]

message-terminator [string]

no-response [yes | no]

edit protocol [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats protocol

reset-stats protocol [[[name] | [glob] | [regex]] ...]

DISPLAY

list protocol

list protocol [[[name] | [glob] | [regex]] ...]

show running-config protocol

show running-config protocol [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show protocol

show protocol [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete protocol [name]

DESCRIPTION

You can use the protocol component to implement a named generic message parser for use with the message routing framework. You can create a protocol component, and then add it to a virtual server. You do this when you want to separate a stream of bytes, from a connection to a peer, into messages for routing. This also enables a set of iRule commands to create, populate, and route messages.

EXAMPLES

```
create protocol my_protocol defaults-from genericmsg
```

Creates a message protocol component named my_protocol using the system defaults.

```
create protocol my_protocol { welcome-message hello }
```

Creates a protocol instance named my_protocol that sends a welcome message of "hello" to any new connection.

OPTIONS

defaults-from

Specifies the protocol that you want to use as the parent protocol. The new protocol inherits all of the settings and values from the specified parent protocol. The default value is genericmsg.

description

User defined description.

disable-parser

When set to yes, the generic message parser is disabled. The parser ignores all incoming packets and does not directly send message data. This mode supports iRule script protocol implementations that generate messages from the incoming transport stream and send messages on the outgoing transport stream.

max-egress-buffer

Specifies the maximum size of the send buffer in bytes. If the number of bytes in the send buffer for a connection exceeds this value, the generic message parser stops receiving outgoing messages from the router until the size of the buffer drops below this setting. The default value is 32768.

max-message-size

Specifies the maximum size of a received message. If a message exceeds this size, the connection is reset. The default value is 32768.

message-terminator

Specifies the string of characters used to terminate a message. If the message-terminator parameter is empty, the generic message parser does not separate the input stream into messages. The default value is \n.

no-response

When set to yes, matching of responses to requests is disabled. The default value is no.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh, ltm message-routing generic route ltm message-routing generic protocol

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2014. All rights reserved.

BIG-IP 2014-10-17 ltm message-routing generic protocol(1)

ltm message-routing generic route

NAME

route - Configures a static route the generic message parser uses to route generic message protocol messages.

MODULE

ltm message-routing generic

SYNTAX

Configure the route component within the ltm message-routing generic module using the syntax shown in the following sections.

CREATE/MODIFY

create route [name]

modify route [name]

options:

app-service [[string] | none]

description [string]

destination-address [string]

peer-selection-mode [sequential | ratio]

peers { [peer-name] }

source-address [string]

edit route [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats route

reset-stats route [[[name] | [glob] | [regex]] ...]

DISPLAY

list route

list route [[[name] | [glob] | [regex]] ...]

DELETE

delete route [name]

DESCRIPTION

You can use the route component to manage a generic message static route.

EXAMPLES

```
create route my_route
```

Creates a static route named my_route that uses a wildcard value for the source-address and destination-address parameters. This acts as a default route.

```
create route my_route { destination-address helpdesk peers add { peer1 peer2 }
```

Creates a static route named my_route that contains two peers, peer1 and peer2. Messages routed with a destination-address of helpdesk are routed to a pool member contained in peer1 or peer2, based on the specified peer-selection-mode.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

destination-address

Specifies the destination address of the route. If this parameter is not present, the generic message parser considers the destination-address as a wildcard that matches all message destination addresses. The default value is none.

description

User defined description.

peer-selection-mode

Specifies the method the generic message parser uses to select a peer from the specified list of peers. The default value is sequential.

sequential

Specifies that the generic message parser selects the first peer in the list of peers. If the protocol retransmits the message, the generic message parser uses another pool member in the first peer. If all pool members in a peer are unavailable, the generic message parser uses the next peer in the list.

ratio

Specifies that the generic message parser selects a peer from a list of peers based on the relative ratio values of each peer. For example if three peers have ratios of 1, 1, and 2, the first 2 peers have a 25% (1/4) probability of being selected and the third peer has a 50% (2/4) probability of being selected.

peers

Specifies a list of peers.

source-address

Specifies the source address of the route. If this parameter is not present, the generic message parser considers the source-address as a wildcard that matches all message sources addresses. The default value is none.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, tmsh, ltm message-routing generic peer ltm message-routing generic router

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2014. All rights reserved.

BIG-IP 2017-03-15 ltm message-routing generic route(1)

ltm message-routing generic router

NAME

router - Configures a message router instance for routing generic message protocol messages.

MODULE

ltm message-routing generic

SYNTAX

Configure the router component within the ltm message-routing generic module using the syntax shown in the following sections.

CREATE/MODIFY

create router [name]

modify router [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

ha-message-sweeper-interval [integer]

ignore-client-port [yes | no]

max-pending-bytes [integer]

max-pending-messages [integer]

max-retries [integer]

mirrored [enabled | disabled]

per-peer-stats [enabled | disabled]

routes { [route-name] }

traffic-group [[string] | default | non-default | none]

use-local-connection [yes | no]

irule-scope-message [yes | no]

edit router [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats router

reset-stats router [[[name] | [glob] | [regex]] ...]

DISPLAY

list router

list router [[[name] | [glob] | [regex]] ...]

show running-config router

show running-config router [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show router

show router [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete router [name]

DESCRIPTION

You can use the router component to manage a generic message router instance. All virtual servers containing

the same router instance share the same route table and can route messages between peers.

EXAMPLES

```
create router my_router defaults-from messengerouter
```

Creates a message router instance named my_router using the system defaults.

```
create router my_router { routes add { route1 route2 } }
```

Creates a router instance named my_router that contains two static routes, route1 and route2.

OPTIONS

app-service

Specifies the name of the application service to which the router belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the router. Only the application service can modify or delete the router.

defaults-from

Specifies the profile that you want to use as the parent profile. The new profile inherits all of the settings and values from the specified parent profile. The default value is messengerouter.

description

User defined description.

ignore-client-port

If set to yes, the system ignores the remote port on clientside connections (connections where the peer connected to the BIG IP system) when searching for an existing connection. The default value is no.

irule-scope-message

If set to yes, iRule events are scoped to the message executing (each message has its own execution context). This allows multiple messages to process iRules concurrently without waiting for commands from other messages. If set to no, all genericmsg iRule events are scoped to the connection flow (that is all share a single execution context per flow) and therefore are forced to execute one at a time even though they are processing independent messages (this is the legacy mode). The default value is no. Some existing iRules might need to be adjusted to use the per-message scope.

ha-message-sweeper-interval

Specifies the maximum time in milliseconds that a message will be held on the standby device as it waits for the active device to route the message. Messages on the standby device held for longer then the configurable sweeper interval, will be dropped. The default value is 1000 milliseconds.

inherited-traffic-group

Read-only property that indicates if the traffic-group is inherited from the parent folder.

max-pending-bytes

Limits the number of bytes contained within messages held pending while waiting for a connection to a peer. If irule-scope-message is set to yes, also limits the number of bytes contained within messages that may concurrently process iRule events. Once reached any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 32768.

max-pending-messages

Limits the number of messages held pending while waiting for a connection to a peer. If irule-scope-message is set to yes, also limits the number of messages that may concurrently process iRule events. Once the limit is reached, any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 64.

max-retries

This attribute sets the maximum number of time a message may be resubmitted for rerouting by the MR::retry iRule command. The default value is 1.

mirrored

If enabled, connection created on the active device of the traffic-group specified, will be mirrored on the standby device. Messages processed on the active device will also be mirrored and perform equivalent processing on the standby device.

per-peer-stats

If enabled, the profile specific statistics will be captured for each pool member. The default value is disabled.

traffic-group

Specifies the traffic group on which the router is active. The default traffic group is inherited from the containing folder.

inherited-traffic-group

Read-only property that indicates if the traffic-group is inherited from the parent folder.

use-local-connection

If true, the router will route a message to an existing connection on the same TMM as the message was received on. If an existing connection is not found, it will route the message through an existing connection based on a deterministic algorithm that may be on another TMM. If a matching existing connection is not found, it will create a connection on the current TMM. Setting this flag may limit the number of connections that are created to a peer.

routes

Specifies a list of static routes for the router instance to use.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh, ltm message-routing generic route ltm message-routing generic protocol

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2014. All rights reserved.

BIG-IP 2019-12-19 ltm message-routing generic router(1)

ltm message-routing generic transport-config

NAME

transport-config - Configures a message transport-config instance for routing generic message protocol messages.

MODULE

ltm message-routing generic

SYNTAX

Configure the transport-config component within the ltm message-routing generic module using the syntax shown in the following sections.

CREATE/MODIFY

```
create transport-config [name]
modify transport-config [name]
options:
  app-service [[string] | none]
  description [string]
  profiles [add | delete | replace-all-with] {
    [profile_name ...] {
context [all | clientside | serverside]  read-only attribute for v12.0.0 or greater.
    }
  }
  rules { [none | [rule_name ... ] ] }
  source-address-translation {
    options:
pool [ [pool_name] | none]
type [ automap | snat | none ]
  }
  source-port [integer]
  source-port-mode [change | preserve | preserve-strict]
```

```
edit transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list transport-config
list transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete transport-config [name]
```

DESCRIPTION

You can use the transport-config component to define the profiles, rules, and source-address-translation of an outgoing connection.

EXAMPLES

```
create transport-config my_transport-config
```

Creates a transport-config instance named my_transport-config using the system defaults.

```
create transport-config my_transport-config { profiles add { my_genericmsg my_tcp } }
```

Creates a transport-config instance named my_transport-config that will use two profiles, my_genericmsg and my_tcp, to create and configure an outgoing connection. The outgoing connection is automatically configured with the router instance that created the connection.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description
User defined description.

profiles
Specifies a list of profiles that the outgoing connection uses to direct and manage traffic. The default value is none.

rules
Specifies a list of iRules, separated by spaces, that customize the transport configuration to direct and manage traffic. The default value is none.

source-address-translation
Specifies the type of source address translation enabled for the transport configuration, as well as the pool that the source address translation uses.

pool Specifies the name of a SNAT pool used by the specified transport configuration.

type Specifies the type of source address translation associated with the specified transport configuration.

The options are:

automap
Specifies the use of self IP addresses for transport configuration server source address translation.

none Specifies no source address translation is used by the transport configuration.

snat Specifies the use of a SNAT pool of translation addresses for virtual server source address translation.

source-port
Specifies the source port to be used for the connection being created. If source-port-mode is change, this setting has no effect. The default value is 0.

source-port-mode
Specifies how the system should select a source port for the outgoing connection. The default value is change.

The options are:

change
Selects an ephemeral source port for the outgoing connection.

preserve
Attempts to use the value of source-port for the outgoing connection, if specified. Otherwise attempts to preserve the source port of the incoming connection.

preserve-strict
Forces the outgoing connection to use the value of source-port, if specified. Otherwise forces the new connection to preserve the source port of the incoming connection. The system will fail to create a new outgoing connection if the specified source port is already in use.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh, ltm message-routing generic route ltm message-routing generic protocol

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2019-10-07 ltm message-routing generic transport-config(1)

ltm message-routing mqtt peer

NAME
peer - Configures a peer for routing MQTT message protocol messages.

MODULE
ltm message-routing mqtt

SYNTAX

Configure the peer component within the ltm message-routing mqtt module using the syntax shown in the following sections.

CREATE/MODIFY

create peer [name]

modify peer [name]

options:

app-service [[string] | none]

description [string]

pool [name]

transport-config [transport-config]

edit peer [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list peer

list peer [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete peer [name]

DESCRIPTION

You can use the peer component to manage a named MQTT peer. A peer specifies the pool for the MQTT router to use as the destination for MQTT routes. You can also use the peer component to specify how many connections the parser creates to a remote host and what transport the parser uses to establish the connection.

EXAMPLES

```
create peer my_peer { pool my_pool transport { type virtual name my_vip } }
```

Creates a MQTT peer named my_peer which uses the settings of my_vip to establish a connection with a pool member from pool my_pool.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

pool Specifies the name of the pool to which the MQTT parser routes messages.

transport-config

Specifies the name of the transport configuration (ltm message-routing mqtt transport-config) the message router uses to create an outgoing connection.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh, ltm message-routing mqtt route ltm message-routing mqtt profile session

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2017. All rights reserved.

BIG-IP 2018-10-20 ltm message-routing mqtt peer(1)

ltm message-routing mqtt profile router

NAME

router - Configures a MQTT Router profile.

MODULE

ltm message-routing mqtt profile

SYNTAX

Configure the router component within the ltm message-routing mqtt profile module using the syntax shown in the following sections.

CREATE/MODIFY
create router [name]
modify router [name]
options:
 app-service [[string] | none]
 defaults-from [[name] | none]
 description [string]
 max-pending-bytes [integer]
 max-payload-pending-bytes [integer]
 max-pending-messages [integer]
 max-retries [integer]
 per-peer-stats [enabled | disabled]
 route [add | default | delete | none | replace-all-with] [route_name]
 traffic-group [[string] | default | non-default | none]
 use-local-connection [disabled | enabled]

edit router [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties

reset-stats router
reset-stats router [[[name] | [glob] | [regex]] ...]

DISPLAY
list router
list router [[[name] | [glob] | [regex]] ...]
show running-config router
show running-config router [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties
 one-line
 partition

show router
show router [[[name] | [glob] | [regex]] ...]
options:
 (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
 field-fmt

DELETE
delete router [name]

DESCRIPTION
You can use the router component to manage a MQTT router profile.

EXAMPLES
create router my_router_profile defaults-from router

Creates a MQTT router profile named my_router_profile using the system defaults.

create router my_router_profile route route1

Creates a MQTT router profile named my_router_profile with a static route.

OPTIONS
 app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

 defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is router.

 description
User defined description.

 glob
Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

 max-pending-bytes
Specifies the maximum number of bytes contained within pending messages that will be held while waiting for a connection to a peer to be created. Once reached any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 32768.

 max-payload-pending-bytes
Specifies the maximum number of payload bytes contained within pending messages that will be held before exerting flow control and eventually closing the TCP window. The default value is 32768.

 max-pending-messages
Specifies the maximum number of pending messages that will be held while waiting for a connection to a peer to be created. Once reached any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 64.

max-retries

This attribute sets the maximum number of time a message may be resubmitted for rerouting by the MR::retry iRule command. The default value is 1.

per-peer-stats

If enabled, the profile specific statistics will be captured for each pool member. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

route

Specifies a static route.

traffic-group

Specifies the traffic group on which the router is active. The default traffic group is inherited from the containing folder.

inherited-traffic-group

Read-only property that indicates if the traffic-group is inherited from the parent folder.

use-local-connection

Enables or disables a preference for local connections established by the ingress TMM over connections established by other TMM's when selecting the egress connection to destination peer.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016-2017. All rights reserved.

BIG-IP 2019-07-02 ltm message-routing mqtt profile router(1)

ltm message-routing mqtt profile session

NAME

session - Configures a MQTT Session profile.

MODULE

ltm message-routing mqtt profile

SYNTAX

Configure the session component within the ltm message-routing mqtt profile module using the syntax shown in the following sections.

CREATE/MODIFY

create session [name]

modify session [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

edit session [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats session

reset-stats session [[[name] | [glob] | [regex]] ...]

DISPLAY

list session

list session [[[name] | [glob] | [regex]] ...]

show running-config session

show running-config session [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show session

show session [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete session [name]

DESCRIPTION

You can use the session component to manage a MQTT session profile.

EXAMPLES

create session my_session_profile defaults-from session

Creates a MQTT session profile named my_session_profile using the system defaults.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is session.

description

User defined description.

clientid-prefix

Specifies the client-id prefix that will be used when sending the CONNECT message to the broker.

keepalive-int

Specifies the keepalive interval that will be used when sending the CONNECT message to the broker.

peered-session-mode

Set this option to peer a client session with server session. When either side terminates the peered session will terminate immediately.

proxy-topic-prefix

Specifies a prefix that will be added to topic of messages that belong to the proxy.

client-will-handling-mode

Specifies the will handling mode to control will action for ungraceful shutdown of the client session.

server-will-handling-mode

Specifies the will handling mode to control will action for ungraceful shutdown of proxy's session with the broker.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

partition

Displays the administrative partition within which the component resides.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016-2017. All rights reserved.

BIG-IP 2018-10-20 ltm message-routing mqtt profile session(1)

NAME

route - Configures a static route for use in MQTT message routing.

MODULE

ltm message-routing mqtt

SYNTAX

Configure the route component within the ltm message-routing mqtt module using the syntax shown in the following sections.

CREATE/MODIFY

create route [name]

modify route [name]

options:

app-service [[string] | none]

description [string]

peer [peer_name]

virtual-server [virtual-server_name]

edit route [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list route

list route [[[name] | [glob] | [regex]] ...]

show running-config route

show running-config route [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show route

show route [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete route [name]

DESCRIPTION

You can use the route component to define the URI's, virtual server, peers and peer selection mode of a message routing MQTT static route.

EXAMPLES

```
create route my_route
```

Creates a route instance named my_route using the system defaults.

```
create route my_route peer peer1
```

Creates a route instance named my_route that will use peer1 for forwarding messages.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

peer Specifies a peer to use for forwarding messages.

virtual-server

Specifies the virtual server on which connections will be routed to this route. If the virtual server is unset, messages originating on any connection may be routed to the route.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, show, tmsh, ltm message-routing mqtt route

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016-2017. All rights reserved.

BIG-IP 2018-10-20 ltm message-routing mqtt route(1)

ltm message-routing mqtt transport-config

NAME

transport-config - Configures a mqtt transport-config instance for routing mqtt message protocol messages.

MODULE

ltm message-routing mqtt

SYNTAX

Configure the transport-config component within the ltm message-routing mqtt module using the syntax shown in the following sections.

CREATE/MODIFY

```
create transport-config [name]
modify transport-config [name]
options:
  app-service [[string] | none]
  description [string]
  profiles [add | delete | replace-all-with] {
    [profile_name ...] {
context [all | clientside | serverside]  read-only attribute for v12.0.0 or greater.
    }
  }
  rules { [none | [rule_name ... ] ] }
  source-address-translation {
    options:
  pool [ [pool_name] | none]
  type [ automap | snat | none ]
  }
  source-port [integer]
  source-port-mode [change | preserve | preserve-strict]
```

```
edit transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list transport-config
list transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete transport-config [name]
```

DESCRIPTION

You can use the transport-config component to define the profiles, rules, and source-address-translation of an outgoing connection.

EXAMPLES

```
create transport-config my_transport-config
```

Creates a transport-config instance named my_transport-config using the system defaults.

```
create transport-config my_transport-config { profiles add { my_mqttmsg my_tcp } }
```

Creates a transport-config instance named my_transport-config that will use two profiles, my_mqttmsg and my_tcp, to create and configure an outgoing connection. The outgoing connection is automatically configured with the router instance that created the connection.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

profiles

Specifies a list of profiles that the outgoing connection uses to use to direct and manage traffic. The default value is none.

rules

Specifies a list of iRules, separated by spaces, that customize the transport configuration to direct and manage traffic. The default value is none.

source-address-translation

Specifies the type of source address translation enabled for the transport configuration, as well as the

pool that the source address translation uses.

pool Specifies the name of a SNAT pool used by the specified transport configuration.

type Specifies the type of source address translation associated with the specified transport configuration.

The options are:

automap

Specifies the use of self IP addresses for transport configuration server source address translation.

none Specifies no source address translation is used by the transport configuration.

snat Specifies the use of a SNAT pool of translation addresses for virtual server source address translation.

source-port

Specifies the source port to be used for the connection being created. If source-port-mode is change, this setting has no effect. The default value is 0.

source-port-mode

Specifies how the system should select a source port for the outgoing connection. The default value is change.

The options are:

change

Selects an ephemeral source port for the outgoing connection.

preserve

Attempts to use the value of source-port for the outgoing connection, if specified. Otherwise attempts to preserve the source port of the incoming connection.

preserve-strict

Forces the outgoing connection to use the value of source-port, if specified. Otherwise forces the new connection to preserve the source port of the incoming connection. The system will fail to create a new outgoing connection if the specified source port is already in use.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh, ltm message-routing mqtt route ltm message-routing mqtt profile session

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2017. All rights reserved.

BIG-IP 2019-10-07 ltm message-routing mqtt transport-config(1)

ltm message-routing sip peer

NAME

peer - Configures a peer for routing SIP message protocol messages.

MODULE

ltm message-routing sip

SYNTAX

Configure the peer component within the ltm message-routing sip module using the syntax shown in the following sections.

CREATE/MODIFY

create peer [name]

modify peer [name]

options:

app-service [[string] | none]

auto-initialization [enabled | disabled]

auto-initialization-interval [integer]

connection-mode [per-peer | per-blade | per-tmm | per-client | per-client-per-blade |

per-client-per-tmm | per-peer-alternate-tmm | per-client-alternate-tmm]

description [string]

number-connections [integer]

pool [name]

ratio [integer]

transport-config [transport-config]

edit peer [[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties

DISPLAY

list peer

list peer [[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line
- partition

DELETE

delete peer [name]

DESCRIPTION

You can use the peer component to manage a named SIP peer. A peer specifies the pool for the SIP router to use as the destination for SIP routes. You can also use the peer component to specify how many connections the parser creates to a remote host and what transport the parser uses to establish the connection.

EXAMPLES

```
create peer my_peer { pool my_pool transport { type virtual name my_vip } }
```

Creates a SIP peer named `my_peer` which uses the settings of `my_vip` to establish a connection with a pool member from pool `my_pool`.

OPTIONS

`app-service`

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

`auto-initialization`

If enabled, the BIGIP will automatically create outbound connections to the active pool members in the specified pool using the configuration of the specified transport-config. For auto-initialization to attempt to create a connection, the peer must be included in a route that is attached to a router instance. For each router instance that the peer is contained in, a connection will be initiated. The auto-initialization logic will verify at a configurable interval if the a connection exists between the BIG-IP and the pool members of the pool. If a connection does not exist, it will attempt to reestablish one. The default is disabled.

`auto-initialization-interval`

Specifies the interval (in milliseconds) that attempts to initiate a connection occur. Valid ranges are from 500ms to 65535ms. The default is 5000ms.

`connection-mode`

Specifies how the number of connections per host is limited. Note a host (specified in the referred pool) may exist more than one peer object, and those peer objects may have different settings for `connection-mode` and `number_connections`. Thus, these settings specify how messages routed through this peer are distributed between a set of connections, not the maximum number of connections to a specified host. The default value is `per-peer`.

`per-peer`

Specifies the number of connections to a remote host.

`per-blade`

Specifies the number of connections to a remote host per blade in the cluster.

`per-tmm`

Specifies the number of connections to a remote host per TMM in the system.

`per-client`

Specifies the number of connections to a remote host per client connection.

`per-client-per-blede`

Specifies the number of connections to a remote host per client connection for each blade.

`per-client-per-tmm`

Specifies the number of connections to a remote host per client connection for each tmm.

`per-peer-alternate-tmm`

Specifies the number of connections to a remote host. Any new connections will be opened on an alternate TMM to spread CPU usage.

`per-client-alternate-tmm`

Specifies the number of connections to a remote host per client connection. Any new connections will be opened on an alternate TMM to spread CPU usage.

`description`

User defined description.

`number-connections`

Specifies the distribution of connections between the BIG-IP system and a remote host. The default value is 1.

pool Specifies the name of the pool to which the SIP parser routes messages.

ratio

Specifies the ratio the SIP router uses to select a peer from a list of peers for the ltm message-routing sip route. The default value is 1.

transport-config

Specifies the name of the transport configuration (ltm message-routing sip transport-config) the message router uses to create an outgoing connection.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh, ltm message-routing sip route ltm message-routing sip profile session

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2019-11-20 ltm message-routing sip peer(1)

ltm message-routing sip profile router

NAME

router - Configures a Session Initiation Protocol (SIP) Router profile.

MODULE

ltm message-routing sip profile

SYNTAX

Configure the router component within the ltm message-routing sip profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create router [name]
```

```
modify router [name]
```

options:

```
app-service [[string] | none]
```

```
concurrent-sessions-per-subscriber [integer]
```

```
defaults-from [[name] | none]
```

```
description [string]
```

```
dialog-establishment-timeout [integer]
```

```
inherited-traffic-group [true | false]
```

```
log-profile [log profile name | none]
```

```
log-publisher [log publisher name | none]
```

```
max-global-registrations [integer]
```

```
max-pending-bytes [integer]
```

```
max-pending-messages [integer]
```

```
max-retries [integer]
```

```
media-proxy {
```

```
  max-media-sessions [integer]
```

```
  media-inactivity-timeout [integer]
```

```
}
```

```
mirror [enabled | disabled]
```

```
nonregistered-subscriber-callout [enabled | disabled]
```

```
nonregistered-subscriber-listener [enabled | disabled]
```

```
per-peer-stats [enabled | disabled]
```

```
registration-timeout [integer]
```

```
operation-mode [load-balancing | application-level-gateway ]
```

```
routes [add | default | delete | none | replace-all-with] {
```

```
  [route_name] ...
```

```
}
```

```
session {
```

```
  transaction-timeout [integer]
```

```
  max-session-timeout [integer]
```

```
}
```

```
traffic-group [[string] | default | non-default | none]
```

```
use-local-connection [disabled | enabled]
```

```
edit router [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
reset-stats router
```

```
reset-stats router [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list router
list router [ [ [name] | [glob] | [regex] ] ... ]
show running-config router
show running-config router [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

show router
show router [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DELETE

```
delete router [name]
```

DESCRIPTION

You can use the router component to manage a SIP router profile.

EXAMPLES

```
create router my_router_profile defaults-from router
```

Creates a SIP router profile named my_router_profile using the system defaults.

```
create router my_router_profile routes add { route1 route2 }
```

Creates a SIP profile named my_router_profile with two static routes.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is router.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

log-publisher

Specify the name of the log publisher which logs translation events. See help sys log-config for more details on the logging sub-system. Use the "sys log-config publisher" component to set up a log publisher.

log-profile

Specify the name of the ALG log profile which controls the logging of ALG . See help ltm alg-log-profile for more details on the logging profile sub-system. Use the "ltm alg-log-profile profile" component to set up a ALG log profile.

max-pending-bytes

Specifies the maximum number of bytes contained within pending messages that will be held while waiting for a connection to a peer to be created. Once reached any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 32768.

max-pending-messages

Specifies the maximum number of pending messages that will be held while waiting for a connection to a peer to be created. Once reached any additional messages to the peer will be flagged as undeliverable and returned to the originator. The default value is 64.

max-retries

This attribute sets the maximum number of time a message may be resubmitted for rerouting by the MR::retry iRule command. The default value is 1.

mirrored

If enabled, connection created on the active device of the traffic-group specified, will be mirrored on the standby device. Messages processed on the active device will also be mirrored and perform equivalent processing on the standby device.

Important: Changing traffic groups, with Connection Mirroring enabled, drops all mirrored connections and loses all persistence data. If you change traffic groups, mirroring must restart.

per-peer-stats

If enabled, the profile specific statistics will be captured for each pool member. The default value is disabled.

operation-mode

Specifies the behavior of the routing instance. The options are:

load-balancing
Messages will be load balanced.

application-level-gateway
The virtual will act as a forwarding virtual, forwarding messages to the server of peer identified as the destination address of the incoming connection. The messages will be inspected so that associated media flows can be created.

partition
Displays the administrative partition within which the component resides.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

routes
Specifies a list of static routes. The ordering of the route entries is insignificant.

session.transaction-timeout
Specifies the maximum time (in seconds) between a request and its response. A provisional response restarts the timer. This may not affect all transactions. The scenarios where the system waits for response (eg. a final response for REGISTER request), are impacted by dropping any persistent data maintained for the request.

session.max-session-timeout
This attribute is valid when the operation-mode is application-level-gateway. Specifies the maximum duration (in seconds) the media for a call remains active. After this period call media is terminated.

traffic-group
Specifies the traffic group on which the router is active. The default traffic group is inherited from the containing folder.

inherited-traffic-group
Read-only property that indicates if the traffic-group is inherited from the parent folder.

media-proxy.max-media-sessions
This attribute is valid when the operation-mode is application-level-gateway. Specifies the maximum number of media sessions that are allowed per call.

media-proxy.media-inactivity-timeout
This attribute is valid when the operation-mode is application-level-gateway. Specifies the maximum duration (in seconds) that a media flow is active with no RTP packets. After this period the media flow is removed. This attribute is applicable only to RTP packets. Inactivity time applicable to RTCP packets is max-session-timeout.

max-global-registrations
This attribute is valid when the operation-mode is application-level-gateway. Specifies the maximum number of registrations allowed. If the limit is reached the registrations are dropped. A default value of 0 indicates the limit is ignored."

nonregistered-subscriber-callout
This attribute is valid when the operation-mode is application-level-gateway with translation. Enables non-registered subscriber to initiate calls. The default value is enabled.

nonregistered-subscriber-listener
This attribute is valid when the operation-mode is application-level-gateway with translation. Enables creation of an ephemeral listener when a non-registered subscriber initiates a calls. If enabled, Non-registered subscribers will be able to receive incoming calls. The default value is enabled.

concurrent-sessions-per-subscriber
This attribute is valid when the operation-mode is application-level-gateway. Specifies the maximum number of concurrent calls allowed per subscriber. It denotes the concurrent From-URI in the call matching the registration key. If the limit exceeds then the call is dropped. Default value of 0 indicates that the limit is ignored."

registration-timeout
This attribute is valid when the operation-mode is application-level-gateway with LSN source-address-translation. Specifies the maximum duration (in seconds) that a registration entry remains active while a response has not been received from the registrar. After this period the registration entry is removed from the registration table, unless the registrar has responded. A default value of 0 means this timeout is ignored and the Expires header value is used.

dialog_establishment_timeout
This attribute is valid when the operation-mode is application-level-gateway. Specifies the timeout (in seconds) that represents the Timer B as per RFC 3261, the INVITE transaction timeout. The dialog-establishment-timeout is used by the Call Table. The default value is 32 seconds.

use-local-connection
Enables or disables a preference for local connections established by the ingress TMM over connections established by other TMM's when selecting the egress connection to destination peer.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016. All rights reserved.

BIG-IP 2019-07-02 ltm message-routing sip profile router(1)

ltm message-routing sip profile session

NAME

session - Configures a Session Initiation Protocol (SIP) Session profile.

MODULE

ltm message-routing sip profile

SYNTAX

Configure the session component within the ltm message-routing sip profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create session [name]
modify session [name]
options:
  honor-route-mode [loose | strict]
  record-route-mode [single | double]
  service-port [integer]
  allow-unknown-methods [disabled | enabled]
  app-service [[string] | none]
  custom-via [[via-header] | none]
  defaults-from [[name] | none]
  description [string]
  do-not-connect-back [disabled | enabled]
  generate-response-on-failure [disabled | enabled]
  honor-via [disabled | enabled]
  insert-record-route-header [disabled | enabled]
  honor-route-header [disabled | enabled]
  insert-via-header [disabled | enabled]
  maintenance-mode [disabled | enabled]
  loop-detection [disabled | enabled]
  loop-detection-mode [Loose | Strict]
  max-forwards-check [disabled | enabled]
  max-msg-header-count [integer]
  max-msg-header-size [integer]
  max-msg-size [integer]
  passthru-mode [disabled | enabled ]
  persistence {
    persist-key [Call-ID | Custom | Src-Addr]
    persist-timeout [integer]
    persist-type [session | none]
  }
  enable-sip-firewall [no | yes]
```

```
edit session [ [ name ] | [ glob ] | [ regex ] ] ... ]
options:
  all-properties
  non-default-properties
```

```
reset-stats session
reset-stats session [ [ name ] | [ glob ] | [ regex ] ] ... ]
```

DISPLAY

```
list session
list session [ [ name ] | [ glob ] | [ regex ] ] ... ]
show running-config session
show running-config session [ [ name ] | [ glob ] | [ regex ] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show session
show session [ [ name ] | [ glob ] | [ regex ] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DELETE

delete session [name]

DESCRIPTION

You can use the session component to manage a SIP session profile.

EXAMPLES

```
create session my_session_profile defaults-from session
```

Creates a SIP session profile named my_session_profile using the system defaults.

```
create session my_session_profile { insert-record-route-header enabled }
```

Creates a SIP profile named my_session_profile with insertion of record-route header in requests which establish a dialog.

OPTIONS

allow-unknown-methods

If enabled, SIP messages with unknown methods will be parsed and routed. The default value is disabled.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

custom-via Specifies the value the system uses for the Sent-by field of the Via header when the Insert Via Header setting is enabled. Note: The value that you enter must include a format of SIP/SIP-version/protocol, followed by a Sent-By value. For example, SIP/2.0/TCP www.siterequest.com:4343 or SIP/2.0/SCTP 10.10.4.32.
defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is session.

description

User defined description.

do-not-connect-back

Enables or disables whether a connection to a request originator is re-established (if it no longer exists) in order to deliver a response. When disabled, responses that cannot be forwarded using an existing connection are dropped.

generate-response-on-failure

Enables or disables sending failure response messages such as 4xx, 5xx and 6xx, when a SIP request is being dropped. Note: Where it is specified "silently" discarded/dropped, no error response is generated. In any case, a dropped message (request/response) is tracked by the Messages Dropped Statistic.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

honor-via

Enables or disables honoring any via which is not inserted by the system for routing the response. This attribute has no effect if the associated siprouter profile's operation mode attribute is set to \"application level gateway\".

insert-record-route-header

Enables or disables the insertion of a record-route header in requests that establish dialog. When enabled, along with the URI, the custom parameters may be added to facilitate the routing of subsequent requests within this call to avoid route lookup. The record route URI is the local-IP and port of flows that are used for forwarding the message. This attribute has no effect if the associated siprouter profile's operation mode attribute is set to \"application level gateway\".

record-route-mode

"Single" mode will insert one Record-Route header into requests and rewrite the header on response as recommended in RFC 3261. "Double" mode will insert two Record-Route headers as recommended in RFC 5658.

honor-route-header

Enables or disables honoring of a route header in requests of an established dialog. When enabled, topmost route from the route list will be inspected and removed if it belongs to BIGIP. Also message will be routed to location mention in the following route in the route list. This attribute has no effect if the associated siprouter profile's operation mode attribute is set to \"application level gateway\".

honor-route-mode

Determines how to insert and handle Record-Route/Route headers in requests. "Loose" mode derives Record-Route and Route processing from RFC 3261, while "Strict" uses RFC 2543.

insert-via-header

Enables or disables insertion of top via. When enabled, custom parameters to help route the response back are inserted, along with sent-by field of via. The source address:port of the flow forwarding the request is filled as value for sent-by field of the via unless a custom via value is specified. The custom parameters inserted to help routing, helps improve performance as it facilitates routing without any lookup. The via is inserted at egress side of the flow, after the SIP_REQUEST_SEND event. This attribute has no effect if the associated siprouter profile's operation mode attribute is set to \"application level gateway\".

maintenance-mode

Enables or disables maintenance mode. When enabled, SIP response \"503 Service Unavailable\" will be sent for incoming SIP request. SIP response will be dropped.

passthru-mode

Enables or disables passthru mode. When enabled, if the first message received in a flow is not a SIP message, the profile will enter passthru mode. The flow will stay in passthru mode for its lifetime. In passthru mode, all data is passed on without modification or validation.

loop-detection

Enables or disables loop-detection checking. When a loop is detected, the request is discarded. An error response is sent, if configured. A request is detected as seen before (forwarded/spiraled/looped) only if self inserted via is found in the message and the value of its branch param plays a key role in detecting loop versus spiral. Hence enabling via insertion becomes a requirement to do loop detection check. This attribute has no effect if the associated siprouter profile's operation mode attribute is set to \"application level gateway\".

loop-detection-mode

Loop detection logic generates a unique hash per transaction based on several attributes including Cseq, To, From, Call-ID, Route, Request-URI and Chassis serial number. Loop-detection-mode value of Loose will skip "To" attribute in the hash calculation.

max-forwards-check

Enables or disables checking on max-forwards. The max-forwards header field serves to limit the number of hops a request can transit on the way to its destination. If 0, the request is discarded. An error response is sent, if configured. This attribute has no effect if the associated siprouter profile's operation mode attribute is set to \"application level gateway\".

max-msg-header-count

Indicates the maximum count of expected SIP message header fields. A message that exceeds this limit is silently discarded.

max-msg-header-size

Indicates the maximum SIP message header size (in bytes). A message which exceeds this size is silently discarded.

max-msg-size

Indicates the maximum number acceptable SIP message size (in bytes). A message which exceeds this size is silently discarded.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

persistence

Configure the persistence settings.

persistence.persist-key

Specifies the method which should be used to extract the key value that is used to persist on. The options are:

Call-ID

Persist based on the "Call-ID" header field value in the message.

Custom

Persist based on the custom key specified using an iRule.

Src-Addr

Persist based on originating IP address in the message.

The default option is Call-ID.

persistence.persist-timeout

Specifies the timeout value of persistence entries in seconds. Upon receiving of the response for the initial SIP Request message, the persistence record is updated with the persist-timeout value.

persistence.persist-type

Specifies the type of the persistence to be used for the specified "persist-key" attribute value. The options are:

session

Persistence is enabled guaranteeing that messages containing a given persistence key will be delivered to the same peer.

none Persistence is disabled.

The default option is session.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

enable-sip-firewall

Indicates whether to enable the Application Firewall Security policy. When enabled, the configured AFM security features will apply to the virtual server(s) using this SIP session profile.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016. All rights reserved.

BIG-IP 2019-10-01 ltm message-routing sip profile session(1)

ltm message-routing sip route

NAME

route - Configures a static route for use in Session Initiation Protocol (SIP) message routing.

MODULE

ltm message-routing sip

SYNTAX

Configure the route component within the ltm message-routing sip module using the syntax shown in the following sections.

CREATE/MODIFY

create route [name]

modify route [name]

options:

app-service [[string] | none]

description [string]

from-uri [string]

peer-selection-mode [ratio | sequential]

peers { [none] | [peer_name ...] }

request-uri [string]

to-uri [string]

virtual-server [virtual-server_name]

edit route [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list route

list route [[[name] | [glob] | [regex]] ...]

show running-config route

show running-config route [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show route

show route [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete route [name]

DESCRIPTION

You can use the route component to define the URI's, virtual server, peers and peer selection mode of a message routing SIP static route.

EXAMPLES

```
create route my_route
```

Creates a route instance named my_route using the system defaults.

```
create route my_route peers { peer1 peer2 }
```

Creates a route instance named my_route that will use two peers for forwarding messages.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

from-uri

Specifies the patterns to be matched against the From field of a SIP message. This URI will be matched as a case insensitive method. It should be in the form of user@domain. The "sip:" prefix should not be present. Any additional modifiers (for example port or transport) should also not be present. It may begin with a wildcard, "*". If empty, it will be treated as if the entire URI was a wildcard (matching all From-URIs).

peer-selection-mode

Specifies the mode of selecting a peer from a list of peers. The options are:

ratio

Peers are selected based on their weights in comparison with other peers.

sequential

Peers are selected in the order listed. All traffic will route the first peer unless all pool members in the peer are marked down.

peers

Specifies an ordered list of peers to use for forwarding messages.

request-uri

Specifies the patterns to be matched against the request-uri field of a SIP message. This URI will be matched as a case insensitive method. It should be in the form of user@domain. The "sip:" prefix should not be present. Any additional modifiers (for example port or transport) should also not be present. It may begin with a wildcard, "*". If empty, it will be treated as if the entire URI was a * wildcard (matching all Request-URIs).

to-uri

Specifies the patterns to be matched against the To field of a SIP message. This URI will be matched as a case insensitive method. It should be in the form of user@domain. The "sip:" prefix should not be present. Any additional modifiers (for example port or transport) should also not be present. It may begin with a wildcard, "*". If empty, it will be treated as if the entire URI was a wildcard (matching all To-URIs).

virtual-server

Specifies the virtual server on which connections will be routed to this route. If the virtual server is unset, messages originating on any connection may be routed to the route.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, show, tmsh, ltm message-routing sip route

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm message-routing sip route(1)

ltm message-routing sip transport-config

NAME

transport-config - Configures a sip transport-config instance for routing sip message protocol messages.

MODULE

ltm message-routing sip

SYNTAX

Configure the transport-config component within the ltm message-routing sip module using the syntax shown in the following sections.

CREATE/MODIFY

```
create transport-config [name]
modify transport-config [name]
options:
  app-service [[string] | none]
  description [string]
  profiles [add | delete | replace-all-with] {
    [profile_name ...] {
context [all | clientside | serverside] read-only attribute for v12.0.0 or greater.
    }
  }
  rules { [none | [rule_name ... ] ] }
  source-address-translation {
    options:
  pool [ [pool_name] | none]
```

```
type [ automap | snat | none ]
}
source-port [integer]
source-port-mode [change | preserve | preserve-strict]
```

```
edit transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
DISPLAY
list transport-config
list transport-config [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
DELETE
delete transport-config [name]
```

DESCRIPTION

You can use the transport-config component to define the profiles, rules, and source-address-translation of an outgoing connection.

EXAMPLES

```
create transport-config my_transport-config
```

Creates a transport-config instance named my_transport-config using the system defaults.

```
create transport-config my_transport-config { profiles add { my_sipmsg my_tcp } }
```

Creates a transport-config instance named my_transport-config that will use two profiles, my_sipmsg and my_tcp, to create and configure an outgoing connection. The outgoing connection is automatically configured with the router instance that created the connection.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

profiles

Specifies a list of profiles that the outgoing connection uses to use to direct and manage traffic. The default value is none.

rules

Specifies a list of iRules, separated by spaces, that customize the transport configuration to direct and manage traffic. The default value is none.

source-address-translation

Specifies the type of source address translation enabled for the transport configuration, as well as the pool that the source address translation uses.

pool Specifies the name of a SNAT pool used by the specified transport configuration.

type Specifies the type of source address translation associated with the specified transport configuration.

The options are:

automap

Specifies the use of self IP addresses for transport configuration server source address translation.

none Specifies no source address translation is used by the transport configuration.

snat Specifies the use of a SNAT pool of translation addresses for virtual server source address translation.

source-port

Specifies the source port to be used for the connection being created. If source-port-mode is change, this setting has no effect. The default value is 0.

source-port-mode

Specifies how the system should select a source port for the outgoing connection. The default value is change.

The options are:

change

Selects an ephemeral source port for the outgoing connection.

preserve
Attempts to use the value of source-port for the outgoing connection, if specified. Otherwise attempts to preserve the source port of the incoming connection.

preserve-strict
Forces the outgoing connection to use the value of source-port, if specified. Otherwise forces the new connection to preserve the source port of the incoming connection. The system will fail to create a new outgoing connection if the specified source port is already in use.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh, ltm message-routing sip route ltm message-routing sip profile session

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2019-10-07 ltm message-routing sip transport-config(1)

ltm monitor diameter

NAME

diameter - Configures a monitor for Diameter protocol resources.

MODULE

ltm monitor

SYNTAX

Configure the diameter component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create diameter [name]

modify diameter [name]

options:

acct-application-id [[integer] | none]

app-service [[string] | none]

auth-application-id [[integer] | none]

defaults-from [name]

description [string]

host-ip-address [[ip address] | none]

interval [integer]

manual-resume [enabled | disabled]

mode [tcp | mr-tcp | mr-sctp]

origin-host [[ip address] | none]

origin-realm [[hostname] | none]

product-name [name]

time-until-up [integer]

timeout [integer]

up-interval [integer]

vendor-id [integer]

vendor-specific-acct-application-id [[integer] | none]

vendor-specific-auth-application-id [[integer] | none]

vendor-specific-vendor-id [[integer] | none]

edit diameter [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list diameter

list diameter [[[name] | [glob] | [regex]] ...]

show diameter [[[name] | [glob] | [regex]] ...]

show running-config diameter

show running-config diameter [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete diameter [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the diameter component to configure a custom monitor, or you can use the default Diameter monitor that the Local Traffic Manager provides. This type of monitor checks the health of Diameter protocol resources.

EXAMPLES

```
create diameter my_diameter defaults-from diameter
```

Creates a monitor named my_diameter that inherits properties from the default Diameter monitor.

```
list diameter
```

Displays the properties of all of the Diameter monitors.

OPTIONS

`acct-application-id`

Specifies the ID of the accounting portion of a Diameter application. If you specify this option, you must also specify a value for the `auth-application-id` option. The default value is none.

Note that the `acct-application-id` and `auth-application-id` attribute-value-pair (AVP), and the `vendor-specific-auth-application-id` and `vendor-specific-acct-application-id` AVP are mutually exclusive. You can only specify one of these AVPs.

`app-service`

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

`auth-application-id`

Specifies the ID of the authentication and authorization portion of a Diameter application. If you specify this option, you must also specify a value for the `acct-application-id` option. The default value is none.

Note that the `acct-application-id` and `auth-application-id` attribute-value-pair (AVP), and the `vendor-specific-auth-application-id` and `vendor-specific-acct-application-id` AVP are mutually exclusive. You can only specify one set of these AVPs.

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `firepass`.

`description`

User defined description.

`glob` Displays the items that match the `glob` expression. See help `glob` for a description of `glob` expression syntax.

`host-ip-address`

Specifies the IP address of the sender of the Diameter message for the Diameter protocol peer discovery feature. The default value is none.

`interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the `up-interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

`manual-resume`

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value is disabled.

Note that if you set the `manual-resume` option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

`mode` Specifies the transport protocol that the monitor uses to communicate with the target. The default mode is `tcp`. The options are:

`tcp` Specifies that the monitor uses TCP to communicate with the target.

`mr-tcp`

Specifies that the monitor uses TCP with the message-routing framework to communicate with the target.

`mr-sctp`

Specifies that the monitor uses SCTP with the message-routing framework to communicate with the target.

The modes beginning with "mr-" use an in-TMM monitor based on the message-routing framework (MRF). The other modes use the legacy external monitor.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`origin-host`

Specifies the IP address from which the Diameter message originates. The default value is the fully-

qualified domain name of the BIG-IP system.

`origin-realm`

Specifies the realm in which the host from which the Diameter message originates resides. The default value is `f5.com`.

`partition`

Displays the administrative partition within which the component resides.

`product-name`

Specifies the vendor-assigned name of the Diameter application. The value of this option must remain constant across firmware revisions for the same product. The default value is `F5 BIGIP Diameter Health Monitoring`.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`time-until-up`

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

`up-interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

`vendor-id`

Specifies the IANA SMI Network Management Private Enterprise Code assigned to the vendor of the Diameter application. The default value is 3375.

`vendor-specific-acct-application-id`

Specifies the ID of the vendor-specific accounting portion of a Diameter application. If you specify this option, you must also specify a value for the `vendor-specific-auth-application-id` option. The default value is none.

Note that the `acct-application-id` and `auth-application-id` attribute-value-pair (AVP), and the `vendor-specific-auth-application-id` and `vendor-specific-acct-application-id` AVP are mutually exclusive. You can only specify one of these AVPs.

`vendor-specific-auth-application-id`

Specifies the ID of the vendor-specific authentication and authorization portion of a Diameter application. If you specify this option, you must also specify a value for the `vendor-specific-acct-application-id` option. The default value is none.

Note that the `acct-application-id` and `auth-application-id` attribute-value-pair (AVP), and the `vendor-specific-auth-application-id` and `vendor-specific-acct-application-id` AVP are mutually exclusive. You can only specify one of these AVPs.

`vendor-specific-vendor-id`

Specifies the ID of a vendor-specific Diameter application. The system uses this ID to advertise support for the application. The default value is none.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012. All rights reserved.

BIG-IP 2019-10-30 itm monitor diameter(1)

NAME

dns - Configures a Domain Name System (DNS) monitor.

MODULE

ltm monitor

SYNTAX

Configure the dns component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create dns [name]

modify dns [name]

options:

accept-rcode [no-error | anything]

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

answer-contains [query-type | any-type | anything]

app-service [[string] | none]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

manual-resume [enabled | disabled]

qname [string]

qtype [a | aaaa]

recv [none | [string]]

reverse [enabled | disabled]

time-until-up [integer]

timeout [integer]

transparent [disabled | enabled]

up-interval [integer]

edit dns [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list dns

list dns [[[name] | [glob] | [regex]] ...]

show dns [[[name] | [glob] | [regex]] ...]

show running-config dns

show running-config dns [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete dns [name]

Note: You cannot delete default monitors.

RUN

run dns [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop dns [name]

DESCRIPTION

You can use the dns component to configure a custom monitor. This type of monitor verifies the DNS service by attempting to send DNS requests generated using the parameters provided to a pool, pool member, or virtual server and validating the DNS response.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create dns my_dns defaults-from dns qname www.test.com
```

Creates a monitor named my_dns that inherits properties other than qname from the default DNS monitor.

```
list dns
```

Displays the properties of all of the DNS monitors.

```
run dns my_dns destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_dns against a target node at 10.10.10.10:80.

stop dns my_dns

Cancels a one-shot test of the custom monitor my_dns in progress.

show dns my_dns test-result

Displays the result of the most recent one-shot test of the custom monitor my_dns.

OPTIONS

accept_rcode

Specifies the RCODE required in the response for an 'up' status. The default value is no-error.

The options are:

no-error

Specifies that the status of the node will be marked up if the received dns message has RCODE = NOERROR.

anything

Specifies that the status of the node will be marked up irrespective of the RCODE in the dns message received.

adaptive

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

adaptive-divergence-type

Specifies whether the adaptive-divergence-value is relative or absolute.

adaptive-divergence-value

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%.) If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool member or node would be marked down.)

adaptive-limit

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

adaptive-sampling-timespan

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

answer_contains

Specifies the record types required in the answer section of the response in order to mark the status of a node up. The default value is query-type.

The options are:

query-type

Specifies that the response should contain at least one answer of which the resource record type matches the qtype.

any-type

Specifies that the dns message should contain at least one answer.

anything

Specifies that an empty answer section is enough to mark the status of the node up.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is dns.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member

and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port you specify, route the check through the IP address and port supplied by the pool member, and mark the pool member (the gateway) up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

qname

Specifies the query name that the monitor send a DNS query for. The default value is Enter a query name.

qtype

Specifies the query type of that the monitor sends a query. The default value is a.

The options are:

a Specifies that the monitor will send a DNS query of type A.

aaaa Specifies that the monitor will send a DNS query of type AAAA.

recv Specifies the ip address that the monitor looks for in the dns response's resource record sections. The ip address should be specified in the dotted-decimal notation or ipv6 notation. The default value is none. If no recv value is specified, then the dns message will be checked against accept_rcode and answer_contains monitor parameters respectively.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reverse

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful receive string match marks the monitored object down instead of up. You can use the this mode only if you configure recv option.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016-2017. All rights reserved.

BIG-IP 2017-08-16 Itm monitor dns(1)

Itm monitor external

NAME

external - Configures an external monitor.

MODULE

itm monitor

SYNTAX

Configure the external component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

create external [name]

modify external [name]

options:

args [[arguments] | none]

app-service [[string] | none]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

manual-resume [enabled | disabled]

run [none | [external monitor]]

time-until-up [integer]

timeout [integer]

user-defined [[name] [value] | [name] none]

up-interval [integer]

edit external [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list external

list external [[[name] | [glob] | [regex]] ...]

show external [[[name] | [glob] | [regex]] ...]

show running-config external

show running-config external [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete external [name]

Note: You cannot delete default monitors.

RUN

run external [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop external [name]

DESCRIPTION

You can use the external component to configure a custom monitor, or you can use the default external monitor that the Local Traffic Manager provides. Using this type of monitor, you can use your own programs to monitor services.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create external my_external defaults-from external
```

Creates a monitor named my_external that inherits properties from the default external monitor.

```
list external
```

Displays the properties of all of the external monitors.

```
run external my_external destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_external against a target node at 10.10.10.10:80.

```
stop external my_external
```

Cancels a one-shot test of the custom monitor my_external in progress.

```
show external my_external test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_external.

OPTIONS

args Specifies any command-line arguments that the external program requires. The default value is none.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is external.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

run Specifies the external monitor file to be executed by the external monitor. The default value is none.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 zero, which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

user-defined

Specifies any user-defined command-line arguments and variables that the external program requires. Use the following syntax to specify a user defined parameter.

modify external my_external user-defined my_param_name my_param_value

Use the following syntax to remove a user defined parameter.

modify external my_external user-defined my_param_name none

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor external(1)

ltm monitor firepass

NAME

firepass - Configures a FirePass(r) monitor.

MODULE

ltm monitor

SYNTAX

Configure the firepass component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create firepass [name]

modify firepass [name]

options:

app-service [[string] | none]

cipherlist [list]

concurrency-limit [integer]

defaults-from [name]

description [string]

```
destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
interval [integer]
max-load-average [integer]
password [password]
time-until-up [integer]
timeout [integer]
up-interval [integer]
username [ [name] | none]
```

```
edit firepass [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
DISPLAY
list firepass
list firepass [ [name] | [glob] | [regex] ] ... ]
show firepass [ [name] | [glob] | [regex] ] ... ]
show running-config firepass
show running-config firepass [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

```
DELETE
delete firepass [name]
```

Note: You cannot delete default monitors.

```
RUN
run firepass [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

```
STOP
stop firepass [name]
```

DESCRIPTION

You can use the firepass component to configure a custom monitor, or you can use the default Firepass monitor that the Local Traffic Manager provides. This type of monitor checks the health of FirePass systems.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create firepass my_firepass defaults-from firepass
```

Creates a monitor named my_firepass that inherits properties from the default Firepass monitor.

```
list firepass
```

Displays the properties of all of the Firepass monitors.

```
run firepass my_firepass destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_firepass against a target node at 10.10.10.10:80.

```
stop firepass my_firepass
```

Cancels a one-shot test of the custom monitor my_firepass in progress.

```
show firepass my_firepass test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_firepass.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

cipherlist

Specifies the list of ciphers for this monitor. The default value is HIGH:!ADH.

concurrency-limit

Specifies the maximum percentage of licensed connections currently in use under which the monitor marks the FirePass system up. The default value is 95.

For example, a value of 95 percent means that the monitor marks the FirePass system up until 95 percent of licensed connections are in use. When the number of in-use licensed connections exceeds 95 percent, the monitor marks the FirePass system down.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is firepass.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

max-load-average

Specifies the number that the monitor uses to mark the FirePass system up or down. The system compares the value of this option to a one-minute average of the FirePass system load. When the FirePass system-load average falls within the specified value, the monitor marks the FirePass system up. When the average exceeds the value, the monitor marks the system down.

The default value is 12.0.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password, if the monitored target requires authentication. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is gtmuser.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor firepass(1)

ltm monitor ftp

NAME

ftp - Configures a File Transfer Protocol (FTP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the ftp component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create ftp [name]

modify ftp [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

filename [[filename] | none]

interval [integer]

manual-resume [enabled | disabled]

mode [passive | port]

password [none | [password]]

time-until-up [integer]

timeout [integer]

up-interval [integer]

username [name]

edit ftp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ftp

list ftp [[[name] | [glob] | [regex]] ...]

show ftp [[[name] | [glob] | [regex]] ...]

show running-config ftp

show running-config ftp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete ftp [name]

Note: You cannot delete default monitors.

RUN

run ftp [name] [destination [[ipv4 address[:port]] | [ipv6 address[.port]]]]

STOP

stop ftp [name]

DESCRIPTION

You can use the ftp component to configure a custom monitor, or you can use the default FTP monitor that the Local Traffic Manager provides. This type of monitor verifies the FTP service by attempting to download a specific file to the /var/tmp directory on the system. Once downloaded successfully, the file is not saved.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create ftp my_ftp defaults-from ftp
```

Creates a monitor named my_ftp that inherits properties from the default FTP monitor.

```
list ftp
```

Displays the properties of all of the FTP monitors.

```
run ftp my_ftp destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_ftp against a target node at 10.10.10.10:80.

```
stop ftp my_ftp
```

Cancels a one-shot test of the custom monitor my_ftp in progress.

```
show ftp my_ftp test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_ftp.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is ftp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the address and port supplied by a pool member.

***:port**

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

filename

Specifies the full path and file name of the file that the system attempts to download. The health check is successful if the system can download the file. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor

checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

mode Specifies the data transfer process (DTP) mode. The default value is passive.

The options are:

passive

Specifies that the monitor sends a data transfer request to the FTP server. When the FTP server receives the request, the FTP server then starts and establishes the data connection.

port Specifies that the monitor starts and establishes the data connection with the FTP server.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

password

Specifies the password, if the monitored target requires authentication. The default value is none.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 zero, which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor ftp(1)

ltm monitor gateway-icmp

NAME

gateway-icmp - Configures a Gateway Internet Control Message Protocol (ICMP) monitor.

MODULE
 itm monitor

SYNTAX
 Configure the gateway-icmp component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create gateway-icmp [name]
modify gateway-icmp [name]
options:
  adaptive [enabled | disabled]
  adaptive-divergence-type [relative | absolute]
  adaptive-divergence-value [integer]
  adaptive-limit [integer]
  adaptive-sampling-timespan [integer]
  app-service [[string] | none]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  interval [integer]
  manual-resume [enabled | disabled]
  time-until-up [integer]
  timeout [integer]
  transparent [enabled | disabled]
  up-interval [integer]
```

```
edit gateway-icmp [ [ name ] | [ glob ] | [ regex ] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list gateway-icmp
list gateway-icmp [ [ name ] | [ glob ] | [ regex ] ] ... ]
show gateway-icmp [ [ name ] | [ glob ] | [ regex ] ] ... ]
show running-config gateway-icmp
show running-config gateway-icmp [ [ name ] | [ glob ] | [ regex ] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

DELETE

```
delete gateway-icmp [name]
```

Note: You cannot delete default monitors.

RUN

```
run gateway-icmp [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

STOP

```
stop gateway-icmp [name]
```

DESCRIPTION

You can use the gateway-icmp component to configure a custom monitor, or you can use the default Gateway ICMP monitor that the Local Traffic Manager provides. This type of monitor monitors a pool that implements gateway fail-safe for high availability.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create gateway-icmp my_icmp defaults-from gateway_icmp
```

Creates a monitor named my_icmp that inherits properties from the default Gateway ICMP monitor.

```
list gateway-icmp
```

Displays the properties of all of the Gateway ICMP monitors.

```
run gateway-icmp my_gateway-icmp destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_gateway-icmp against a target node at 10.10.10.10:80.

```
stop gateway-icmp my_gateway-icmp
```

Cancels a one-shot test of the custom monitor my_gateway-icmp in progress.

```
show gateway-icmp my_gateway-icmp test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_gateway-icmp.

OPTIONS

adaptive

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

adaptive-divergence-type

Specifies whether the adaptive-divergence-value is relative or absolute.

adaptive-divergence-value

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%.) If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool member or node would be marked down.)

adaptive-limit

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

adaptive-sampling-timespan

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is gateway_icmp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool member (the gateway) is marked up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor gateway-icmp(1)

ltm monitor http

NAME

http - Configures a Hypertext Transfer Protocol (HTTP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the http component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create http [name]

modify http [name]

options:

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

app-service [[string] | none]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

ip-tos [integer]

manual-resume [enabled | disabled]

password [none | [password]]

recv [none | [string]]

recv-disable [none | [string]]

reverse [enabled | disabled]
ip-dscp [integer]
send [none | [string]]
time-until-up [integer]
timeout [integer]
transparent [enabled | disabled]
up-interval [integer]
username [[name] | none]

edit http [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list http

list http [[[name] | [glob] | [regex]] ...]

show http [[[name] | [glob] | [regex]] ...]

show running-config http

show running-config http [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition
test-result

DELETE

delete http [name]

Note: You cannot delete default monitors.

RUN

run http [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop http [name]

DESCRIPTION

You can use the http component to configure a custom monitor, or you can use the default HTTP monitor that the Local Traffic Manager provides. This type of monitor verifies the HTTP service by attempting to receive specific content from a Web page.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create http my_http defaults-from http
```

Creates a monitor named my_http that inherits properties from the default HTTP monitor.

```
list http
```

Displays the properties of all of the HTTP monitors.

```
run http my_http destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_http against a target node at 10.10.10.10:80.

```
stop http my_http
```

Cancels a one-shot test of the custom monitor my_http in progress.

```
show http my_http test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_http.

OPTIONS

adaptive

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

adaptive-divergence-type

Specifies whether the adaptive-divergence-value is relative or absolute.

adaptive-divergence-value

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%.) If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool

member or node would be marked down.)

adaptive-limit

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

adaptive-sampling-timespan

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is http.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *:*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool member (the gateway) is marked up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

ip-dscp

Specifies the differentiated services code point (DSCP). DSCP is a 6-bit value in the Differentiated Services (DS) field of the IP header. It can be used to specify the quality of service desired for the packet. The valid range for this value is 0 to 63 (hex 0x0 to 0x3f). The default value is zero.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-disable

Specifies a text string that the monitor looks for in the returned resource. If the text string is matched in the returned resource, the corresponding node or pool member is marked session disabled. The default value is none.

You specify a `recv-disable` string in the same way that you specify a `recv` string.

If you specify a `recv-disable` string, you must also specify a `recv` string. You cannot specify a `recv-disable` string, if the `reverse` option is enabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

reverse

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use this mode only if you configure both the `send` and `recv` options.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

`send` Specifies the text string that the monitor sends to the target object.

The default setting is `GET /`, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example, `GET /www/company/index.html`.

Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the `recv` option and ignores the option even if not null.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a `RESET` packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the `interval` option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the `interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `run`, `show`, `stop`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor http(1)

NAME

http2 - Configures a Hypertext Transfer Protocol Version 2 (HTTP/2) monitor.

MODULE

ltm monitor

SYNTAX

Configure the http2 component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create http2 [name]

modify http2 [name]

options:

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

app-service [[string] | none]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

ip-dscp [integer]

manual-resume [enabled | disabled]

password [none | [password]]

recv [none | [string]]

recv-disable [none | [string]]

reverse [enabled | disabled]

send [none | [string]]

ssl-profile [[ssl server profile] | none]

time-until-up [integer]

timeout [integer]

transparent [enabled | disabled]

up-interval [integer]

username [[name] | none]

edit http2 [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list http2

list http2 [[[name] | [glob] | [regex]] ...]

show http2 [[[name] | [glob] | [regex]] ...]

show running-config http2

show running-config http2 [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete http2 [name]

Note: You cannot delete default monitors.

RUN

run http2 [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop http2 [name]

DESCRIPTION

You can use the http2 component to configure a custom monitor, or you can use the default HTTP/2 monitor that the Local Traffic Manager provides. This type of monitor verifies the HTTP/2 service by attempting to receive specific content from a Web page over HTTP/2.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create http2 my_http2 defaults-from http2
```

Creates a monitor named my_http2 that inherits properties from the default HTTP/2 monitor.

```
list http2
```

Displays the properties of all of the HTTP/2 monitors.

run http2 my_http2 destination 10.10.10.10:443

Runs a one-shot test of the custom monitor my_http2 against a target node at 10.10.10.10:443.

stop http2 my_http2

Cancels a one-shot test of the custom monitor my_http2 in progress.

show http2 my_http2 test-result

Displays the result of the most recent one-shot test of the custom monitor my_http2.

OPTIONS

`adaptive`

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

`adaptive-divergence-type`

Specifies whether the adaptive-divergence-value is relative or absolute.

`adaptive-divergence-value`

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%.) If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool member or node would be marked down.)

`adaptive-limit`

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

`adaptive-sampling-timespan`

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

`app-service`

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is http2.

`description`

User defined description.

`destination`

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *:*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool member (the gateway) is marked up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

ip-dscp

Specifies the differentiated services code point (DSCP). DSCP is a 6-bit value in the Differentiated Services (DS) field of the IP header. It can be used to specify the quality of service desired for the packet. The valid range for this value is 0 to 63 (hex 0x0 to 0x3f). The default value is zero.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-disable

Specifies a text string that the monitor looks for in the returned resource. If the text string is matched in the returned resource, the corresponding node or pool member is marked session disabled. The default value is none.

You specify a recv-disable string in the same way that you specify a recv string.

If you specify a recv-disable string, you must also specify a recv string. You cannot specify a recv-disable string, if the reverse option is enabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reverse

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use this mode only if you configure both the send and recv options.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

send Specifies the text string that the monitor sends to the target object.

The default setting is GET /, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html.

Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

ssl-profile

Specifies the server side SSL profile that this monitor will use to ping the monitored node or target.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2019-05-02 ltm monitor http2(1)

ltm monitor https

NAME

https - Configures a Hypertext Transfer Protocol over Secure Socket Layer (HTTPS) monitor.

MODULE

ltm monitor

SYNTAX

Configure the https component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create https [name]

modify https [name]

options:

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

app-service [[string] | none]

cert [[cert list] | none]

cipherlist [string]

compatibility [enabled | disabled]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

interval [integer]

ip-dscp [integer]

key [[key] | none]

manual-resume [enabled | disabled]

password [none | [password]]

recv [none | [string]]

recv-disable [none | [string]]

reverse [enabled | disabled]

send [none | [string]]

ssl-profile [[ssl server profile] | none]

time-until-up [integer]

timeout [integer]

transparent [enabled | disabled]

up-interval [integer]

username [[name] | none]

edit https [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list https

list https [[[name] | [glob] | [regex]] ...]

show https [[[name] | [glob] | [regex]] ...]

show running-config https

show running-config https [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties
one-line
partition
test-result

DELETE

delete https [name]

Note: You cannot delete default monitors.

RUN

run https [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop https [name]

DESCRIPTION

You can use the https component to configure a custom monitor, or you can use the default HTTPS monitor that the Local Traffic Manager provides. This type of monitor verifies the HTTPS service by attempting to receive specific content from a Web page protected by Secure Socket Layer (SSL) security.

Note that one of the pre-configured HTTPS monitors is named https_443, which performs a health check on a server using the IP address supplied by the pool member and port 443.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create https my_https defaults-from https
```

Creates a monitor named my_https that inherits properties from the default HTTPS monitor.

```
list https
```

Displays the properties of all of the HTTPS monitors.

```
run https my_https destination 10.10.10.10:443
```

Runs a one-shot test of the custom monitor my_https against a target node at 10.10.10.10:443.

```
stop https my_https
```

Cancels a one-shot test of the custom monitor my_https in progress.

```
show https my_https test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_https.

OPTIONS

adaptive

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

adaptive-divergence-type

Specifies whether the adaptive-divergence-value is relative or absolute.

adaptive-divergence-value

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%). If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool member or node would be marked down.)

adaptive-limit

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

adaptive-sampling-timespan

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

cert Deprecated since v13.1.0. Use ssl-profile instead. Specifies a file object for a client certificate that the monitor sends to the target SSL server. The default value is none.

cipherlist

Deprecated since v13.1.0. Use ssl-profile instead.

compatibility

Deprecated since v13.1.0. Use ssl-profile instead. Specifies, when enabled, that the SSL options setting (in OpenSSL) is set to ALL. The default value is enabled.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is https.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool member (the gateway) is marked up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

ip-dscp

Specifies the differentiated services code point (DSCP). DSCP is a 6-bit value in the Differentiated Services (DS) field of the IP header. It can be used to specify the quality of service desired for the packet. The valid range for this value is 0 to 63 (hex 0x0 to 0x3f). The default value is zero.

key Deprecated since v13.1.0. Use ssl-profile instead. Specifies the RSA private key if the monitored target requires authentication. The key must be surrounded by quotation marks, for example, key \"client.key\". Note that if you specify a key, you must also specify a value for the cert option. The default value is none.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-disable

Specifies a text string that the monitor looks for in the returned resource. If the text string is matched in the returned resource, the corresponding node or pool member is marked session disabled. The default value is none.

You specify a `recv-disable` string in the same way that you specify a `recv` string.

If you specify a `recv-disable` string, you must also specify a `recv` string. You cannot specify a `recv-disable` string, if the `reverse` option is enabled.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`reverse`

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use this mode only if you configure both the `send` and `recv` options.

The default value is `disabled`, which specifies that the monitor does not operate in reverse mode. The `enabled` value specifies that the monitor operates in reverse mode.

`send` Specifies the text string that the monitor sends to the target object.

The default setting is `GET /`, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example, `GET /www/company/index.html`.

Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the `recv` option and ignores the option even if not null.

`ssl-profile`

Specifies the server side SSL profile that this monitor will use to ping the monitored node or target.

`test-result`

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

`time-until-up`

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a `RESET` packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

`transparent`

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is `disabled`.

`up-interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the `interval` option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the `interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

`username`

Specifies the username, if the monitored target requires authentication. The default value is `none`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `run`, `show`, `stop`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2019-05-02 Itm monitor https(1)

NAME

icmp - Configures an Internet Control Message Protocol (ICMP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the icmp component within the ltm monitor module using the syntax shown in the following sections.

CREATE/MODIFY

create icmp [name]

modify icmp [name]

options:

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

app-service [[string] | none]

defaults-from [name]

description [string]

destination [ip address]

interval [integer]

manual-resume [enabled | disabled]

time-until-up [integer]

timeout [integer]

transparent [enabled | disabled]

up-interval [integer]

edit icmp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list icmp

list icmp [[[name] | [glob] | [regex]] ...]

show icmp [[[name] | [glob] | [regex]] ...]

show running-config icmp

show running-config icmp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete icmp [name]

Note: You cannot delete default monitors.

RUN

run icmp [name] [destination [ip address]]

STOP

stop icmp [name]

DESCRIPTION

You can use the icmp component to configure a custom monitor, or you can use the default ICMP monitor that the Local Traffic Manager provides.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create icmp my_icmp defaults-from icmp
```

Creates a monitor named my_icmp that inherits properties from the default ICMP monitor.

```
list icmp
```

Displays the properties of all of the ICMP monitors.

```
run icmp my_icmp destination 10.10.10.10
```

Runs a one-shot test of the custom monitor my_icmp against a target node at IP address 10.10.10.10.

```
stop icmp my_icmp
```

Cancels a one-shot test of the custom monitor my_icmp in progress.

```
show icmp my_icmp test-result
```

Displays the result of the most recent one-shot test of the custom monitor `my_icmp`.

OPTIONS

`adaptive`

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

`adaptive-divergence-type`

Specifies whether the `adaptive-divergence-value` is relative or absolute.

`adaptive-divergence-value`

Specifies how far from mean latency each monitor probe is allowed to be. If `adaptive-divergence-type` is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%.) If `adaptive-divergence-type` is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool member or node would be marked down.)

`adaptive-limit`

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

`adaptive-sampling-timespan`

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

`app-service`

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `icmp`.

`description`

User defined description.

`destination`

Specifies the IP address of the resource that is the destination of this monitor. The default value is `*`.

Possible values are:

* Specifies to perform a health check on the IP address of the node.

IP address

Specifies to perform a health check on the IP address that you specify, and mark the associated node up or down accordingly.

IP address (with the transparent option enabled)

Specifies to perform a health check on the IP address that you specify, route the check through the IP address of the associated node, and mark the IP address of the associated node up or down accordingly.

This option is required for the command `run`, unless an IP address is specified in the `destination` option for the custom monitor.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the `up-interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

`manual-resume`

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the `manual-resume` option is disabled.

Note that if you set the `manual-resume` option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, `modify`, `run` and `stop`.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-04-05 ltm monitor icmp(1)

ltm monitor imap

NAME

imap - Configures an Internet Message Access Protocol (IMAP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the imap component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create imap [name]

modify imap [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

folder [[name] | none]

interval [integer]

manual-resume [enabled | disabled]

password [none | [password]]

time-until-up [integer]

timeout [integer]

up-interval [integer]

username [[name] | none]

edit imap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY
list imap
list imap [[[name] | [glob] | [regex]] ...]
show imap [[[name] | [glob] | [regex]] ...]
show running-config imap
show running-config imap [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties
 one-line
 partition
 test-result

DELETE
delete imap [name]

Note: You cannot delete default monitors.

RUN
run imap [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP
stop imap [name]

DESCRIPTION

You can use the imap component to configure a custom monitor, or you can use the default IMAP monitor that the Local Traffic Manager provides. This type of monitor verifies IMAP by attempting to open a specified mail folder on a server. This monitor is similar to the POP3 monitor.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create imap my_imap defaults-from imap
```

Creates a monitor named my_imap that inherits properties from the default IMAP monitor.

```
list imap
```

Displays the properties of all of the IMAP monitors.

```
run imap my_imap destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_imap against a target node at 10.10.10.10:80.

```
stop imap my_imap
```

Cancels a one-shot test of the custom monitor my_imap in progress.

```
show imap my_imap test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_imap.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is imap.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

folder

Specifies the name of the folder on the IMAP server that the monitor tries to open. The default value is INBOX.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name

Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor imap(1)

ltm monitor inband

NAME

inband - Configures an Inband (passive) monitor.

MODULE

ltm monitor

SYNTAX

Configure the inband component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create inband [name]
modify inband [name]
options:
  app-service [[string] | none]
  defaults-from [name]
  description [string]
  failure-interval [integer]
  failures [integer]
  response-time [integer]
  retry-time [integer]
```

```
edit inband [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list inband
list inband [ [ [name] | [glob] | [regex] ] ... ]
show inband [ [ [name] | [glob] | [regex] ] ... ]
show running-config inband
show running-config inband [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete inband [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the inband component to configure a custom monitor, or you can use the default Inband monitor that the Local Traffic Manager provides. With this type of monitor the BIG-IP(r) system can perform passive monitoring as part of client requests.

EXAMPLES

```
create inband my_inband defaults-from inband
```

Creates a monitor named my_inband that inherits properties from the default Inband monitor.

```
list inband
```

Displays the properties of all of the Inband monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is inband.

description

User defined description.

failure-interval

Specifies an interval, in seconds. If the number of failures specified in the failures option occurs within this interval, the system marks the pool member as being unavailable. The default value is 30.

failures

Specifies the number of failures that the system allows to occur, within the time period specified in the failure-interval option, before marking a pool member unavailable. The default value is 3, which means that the system marks the pool member unavailable at the fourth failure. The multiple tmm processes use a per-process number to calculate failures, depending on the specified load. For example, for the Round Robin load balancing method, if there are N tmm processes and M pool members, and the Failures setting is set to L, then up to $N*M*L+1$ failures can occur before the system marks the node as down.

Specifying a value of 0 (zero) disables this option.

A failure can be either a failure to connect or a failure of the pool member to respond within the time specified in the response-time option.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

response-time

Specifies an amount of time, in seconds. If the pool member does not respond with data after the specified amount of time has passed, the number of failures in this interval increments by 1. Specifying a value of 0 (zero) disables this option.

retry-time

Specifies the amount of time in seconds after the pool member has been marked unavailable before the system retries to connect to the pool member. Specifying a value of 0 (zero) disables this option.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015. All rights reserved.

BIG-IP 2017-04-05 Itm monitor inband(1)

Itm monitor ldap

NAME

ldap - Configures a Lightweight Directory Access Protocol (LDAP) monitor.

MODULE

itm monitor

SYNTAX

Configure the ldap component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

create ldap [name]

modify ldap [name]

options:

app-service [[string] | none]

base [none | [string]]

chase-referrals [no | yes]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

filter [[LDAP key] | none]

interval [integer]

mandatory-attributes [no | yes]

manual-resume [enabled | disabled]

password [none | [password]]

security [none | ssl | tls]

time-until-up [integer]
timeout [integer]
up-interval [integer]
username [[name] | none]

edit ldap [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties

DISPLAY
list ldap
list ldap [[[name] | [glob] | [regex]] ...]
show ldap [[[name] | [glob] | [regex]] ...]
show running-config ldap
show running-config ldap [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties
 one-line
 partition
 test-result

DELETE
delete ldap [name]

Note: You cannot delete default monitors.

RUN
run ldap [name] [destination [[ipv4 address[:port]] | [ipv6 address[.port]]]]

STOP
stop ldap [name]

DESCRIPTION

You can use the ldap component to configure a custom monitor, or you can use the default LDAP monitor that the Local Traffic Manager provides. This type of monitor verifies the LDAP service by attempting to authenticate the specified user.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create ldap my_ldap defaults-from ldap
```

Creates a monitor named my_ldap that inherits properties from the default LDAP monitor.

```
list ldap
```

Displays the properties of all of the LDAP monitors.

```
run ldap my_ldap destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_ldap against a target node at 10.10.10.10:80.

```
stop ldap my_ldap
```

Cancels a one-shot test of the custom monitor my_ldap in progress.

```
show ldap my_ldap test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_ldap.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

base Specifies the location in the LDAP tree from which the monitor starts the health check. A sample value is dc=bigip-test,dc=net. The default value is none.

chase-referrals

Specifies whether the monitor upon receipt of an LDAP referral entry chases that referral. The default value is yes.

The options are:

no Specifies that the system will treat a referral entry as a normal entry and refrain from querying the remote LDAP server(s) pointed to by the referral entry.

yes Specifies that the system upon receiving any referral entry from the monitored LDAP server query, the system will then query the corresponding LDAP server(s) pointed to by the LDAP query. If the

query for the referral is unsuccessful the system will mark the monitored LDAP server down.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is ldap.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

filter

Specifies an LDAP key for which the monitor searches. A sample value is `objectclass=*`. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

mandatory-attributes

Specifies whether the target must include attributes in its response to be considered up. The default value is no.

The options are:

no Specifies that the system performs only a one-level search (based on the value of the filter option), and does not require that the target returns any attributes.

yes Specifies that the system performs a sub-tree search, and if the target returns no attributes, the target is considered down.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help regex for

a description of regular expression syntax.

security

Specifies the secure communications protocol that the monitor uses to communicate with the target. The default value is none. The options are:

none Specifies that the system does not use a security protocol for communications with the target.

ssl Specifies that the system uses the SSL protocol for communications with the target.

tls Specifies that the system uses the TLS protocol for communications with the target.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor ldap(1)

ltm monitor module-score

NAME

module-score - Configures a Module Score monitor that monitors the performance of a pool or node, rather than the health of the pool or node.

MODULE

ltm monitor

SYNTAX

Configure the module-score component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create module-score [name]

modify module-score [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

interval [integer]

pool [name]

snmp-community [none | [string]]

snmp-ip-address [[ip address] | none]

snmp-port [port]

snmp-version [string]

time-until-up [integer]
timeout [integer]
up-interval [integer]

edit module-score [[[name] | [glob] | [regex]] ...]

options:
all-properties
non-default-properties

DISPLAY

list module-score

list module-score [[[name] | [glob] | [regex]] ...]

show module-score [[[name] | [glob] | [regex]] ...]

show running-config module-score

show running-config module-score [[[name] | [glob] | [regex]] ...]

options:
all-properties
non-default-properties
one-line
partition

DELETE

delete module-score [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the module-score component to configure a custom monitor, or you can use the default Module Score monitor that the Local Traffic Manager provides. This type of monitor enables global and local traffic management systems to load balance in a proportional manner to local traffic management virtual servers associated with the Web Accelerator(tm) and Application Security Manager modules. When you configure a Module Score type of monitor, the local traffic management system uses SNMP to pull the gtm_score values from the downstream virtual servers and set the dynamic ratios on the associated upstream local traffic management pool members or nodes.

More specifically, the Module Score monitor retrieves the gtm_score values from the virtual server and the gtm_vs_score values associated with the virtual server. Then, if a pool name is not specified, this monitor sets the dynamic ratio on the node that is associated with the virtual server.

The BIG-IP(r) system uses the lowest non-zero value of the gtm_vs_score values to set the dynamic ratio. If all gtm_vs_score values are zero, then the gtm_score value is used to set the dynamic ratios. If you specify a pool name in the monitor definition, then the dynamic ratio is set on the pool member.

Note: If you want to distribute traffic to a cluster of WebAccelerator or Application Security Manager virtual servers, you must create a separate custom Module Score monitor for each back-end Local Traffic Manager system.

EXAMPLES

```
create module-score my_module-score defaults-from module_score
```

Creates a monitor named my_module-score that inherits properties from the default Module Score monitor.

```
list module-score
```

Displays the properties of all of the Module Score monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is module_score.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the

resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

pool Specifies a Local Traffic Manager pool name. Use this option if you want the system to set dynamic ratios on a pool member instead of on the associated node for the pool member. The default value is none.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

snmp-community
Specifies the identifier for the SNMP community. The default value is public.

snmp-ip-address
Specifies the IP address of the SNMP server. The default value is none.

snmp-port
Specifies the port associated with the SNMP server. The default value is 161.

snmp-version
Specifies the SNMP version in use by the system. The default value is v2c.

time-until-up
Specifies the amount of time in seconds after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 30 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval
Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014. All rights reserved.

BIG-IP 2017-04-05 Itm monitor module-score(1)

Itm monitor mqtt

NAME
mqtt - Configures a Message Queuing Telemetry Transport (MQTT) monitor.

MODULE
itm monitor

SYNTAX
Configure the mqtt component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY
create mqtt [name]

modify mqtt [name]

options:

app-service [[string] | none]
clientid [[name] | none]
defaults-from [name]
description [string]
destination [ip address][port]
interval [integer]
manual-resume [enabled | disabled]
mqtt-version [[string] | none]
password [none | [password]]
time-until-up [integer]
timeout [integer]
up-interval [integer]
username [[name] | none]

edit mqtt [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list mqtt

list mqtt [[[name] | [glob] | [regex]] ...]

show mqtt [[[name] | [glob] | [regex]] ...]

show running-config mqtt

show running-config mqtt [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

DELETE

delete mqtt [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the mqtt component to configure a custom monitor, or you can use the default mqtt monitor that the Local Traffic Manager provides. This type of monitor verifies the mqtt service by attempting to establish mqtt connection with mqtt server.

EXAMPLES

```
create mqtt my_mqtt defaults-from mqtt
```

Creates a monitor named my_mqtt that inherits properties from the default mqtt monitor.

```
list mqtt
```

Displays the properties of all of the mqtt monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

clientid

Specifies the client identifier to send to MQTT server. The default value is empty.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is mqtt.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port specified in the monitor, routing the check through the IP address and port supplied by the pool member. The pool

member (the gateway) is marked up or down accordingly.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval
Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume
Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

mqtt-version
Specifies a version to communicate with MQTT server. Default value is 3.1.1.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

password
Specifies the password if the monitored target requires authentication. The default value is none.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

time-until-up
Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval
Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username
Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO
create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-09-05 itm monitor mqtt(1)

Itm monitor mssql

NAME
mssql - Configures a Microsoft(r) Windows(r) Structured Query Language (MSSQL) monitor.

MODULE
 ltm monitor

SYNTAX
 Configure the mssql component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create mssql [name]
modify mssql [name]
options:
  app-service [[string] | none]
  count [integer]
  database [ [name] | none]
  debug [no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
  interval [integer]
  manual-resume [enabled | disabled]
  password [none | [password] ]
  recv [none | [string] ]
  recv-column [none | [string] ]
  recv-row [none | [string] ]
  send [none | [string] ]
  time-until-up [integer]
  timeout [integer]
  up-interval [integer]
  username [[name] | none]
```

```
edit mssql [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list mssql
list mssql [ [ [name] | [glob] | [regex] ] ... ]
show mssql [ [ [name] | [glob] | [regex] ] ... ]
show running-config mssql
show running-config mssql [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

DELETE

```
delete mssql [name]
```

Note: You cannot delete default monitors.

RUN

```
run mssql [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

STOP

```
stop mssql [name]
```

DESCRIPTION

You can use the mssql component to configure a custom monitor, or you can use the default Microsoft Windows SQL monitor that the Local Traffic Manager provides. This type of monitor verifies Microsoft Windows SQL-based services.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create mssql my_mssql defaults-from mssql
```

Creates a monitor named my_mssql that inherits properties from the default MSSQL monitor.

```
list mssql
```

Displays the properties of all of the MSSQL monitors.

```
run mssql my_mssql destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_mssql against a target node at 10.10.10.10:80.

```
stop mssql my_mssql
```

Cancels a one-shot test of the custom monitor my_mssql in progress.

show mssql my_mssql test-result

Displays the result of the most recent one-shot test of the custom monitor my_mssql.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

count

Specifies the number of monitor probes after which the connection to the database will be terminated. Count value of zero indicates that the connection will never be terminated. The default value is zero.

database

Specifies the name of the database with which the monitor attempts to communicate. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is mssql.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 30 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database.

If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-column

Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

recv-row

Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

send Specifies the SQL query that the monitor sends to the target database, for example, SELECT count(*) FROM mytable.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 Itm monitor mssql(1)

Itm monitor mysql

NAME

mysql - Configures a MySQL(r) monitor.

MODULE

itm monitor

SYNTAX

Configure the mysql component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

create mysql [name]

modify mysql [name]

options:

app-service [[string] | none]

count [integer]
database [[name] | none]
debug [no | yes]
defaults-from [name]
description [string]
destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
interval [integer]
manual-resume [enabled | disabled]
password [none | [password]]
recv [none | [string]]
recv-column [none | [string]]
recv-row [none | [string]]
send [none | [string]]
time-until-up [integer]
timeout [integer]
up-interval [integer]
username [[name] | none]

edit mysql [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list mysql

list mysql [[[name] | [glob] | [regex]] ...]

show mysql [[[name] | [glob] | [regex]] ...]

show running-config mysql

show running-config mysql [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition
test-result

DELETE

delete mysql [name]

Note: You cannot delete default monitors.

RUN

run mysql [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop mysql [name]

DESCRIPTION

You can use the mysql component to configure a custom monitor, or you can use the default MySQL monitor that the Local Traffic Manager provides. This type of monitor verifies MySQL-based services.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create mysql my_mysql defaults-from mysql
```

Creates a monitor named my_mysql that inherits properties from the default MySQL monitor.

```
list mysql
```

Displays the properties of all of the MySQL monitors.

```
run mysql my_mysql destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_mysql against a target node at 10.10.10.10:80.

```
stop mysql my_mysql
```

Cancels a one-shot test of the custom monitor my_mysql in progress.

```
show mysql my_mysql test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_mysql.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

count

Specifies the number of monitor probes after which the connection to the database will be terminated.

Count value of zero indicates that the connection will never be terminated. The default value is zero.

database

Specifies the name of the database with which the monitor attempts to communicate. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `mysql`.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command `run`, unless an IP address and service port are specified in the `destination` option for the specified custom monitor.

glob Displays the items that match the glob expression. See `help glob` for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 30 seconds.

Important: F5 Networks recommends that when you configure this option and the `up-interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the `manual-resume` option is disabled.

Note that if you set the `manual-resume` option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands `create`, `delete`, `modify`, `run` and `stop`.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database. If you do not specify a value for both the `send` and `recv` options, the monitor performs a simple service check and connect only.

recv-column

Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the `send` and `recv` options. The default value is none.

recv-row

Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the `send` and `recv` options. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`send` Specifies the SQL query that the monitor sends to the target database, for example, `SELECT count(*) FROM mytable`.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the `recv` option and ignores the option even if not null.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `run`, `show`, `stop`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor mysql(1)

ltm monitor nntp

NAME

`nntp` - Configures a Network News Transfer Protocol (NNTP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the `nntp` component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create nntp [name]
```

```
modify nntp [name]
```

options:

```
app-service [[string] | none]
```

```
debug [no | yes]
```

```
defaults-from [name]
```

```
description [string]
```

```
destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ]
```

```
interval [integer]
```

```
manual-resume [enabled | disabled]
```

```
newsgroup [ [name] | none]
```

```
password [none | [password] ]
```

```
time-until-up [integer]
```

```
timeout [integer]
```

```
up-interval [integer]
```

```
username [[name] | none]
```

```
edit nntp [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

DISPLAY
list nntp
list nntp [ [name] | [glob] | [regex] ] ... ]
show nntp [ [name] | [glob] | [regex] ] ... ]
show running-config nntp
show running-config nntp [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result

DELETE
delete nntp [name]
```

Note: You cannot delete default monitors.

```
RUN
run nntp [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]

STOP
stop nntp [name]
```

DESCRIPTION

You can use the nntp component to configure a custom monitor, or you can use the default NNTP monitor that the Local Traffic Manager provides. This type of monitor verifies the Usenet News protocol service by attempting to retrieve a newsgroup identification string from the server.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create nntp my_nntp defaults-from nntp

Creates a monitor named my_nntp that inherits properties from the default NNTP monitor.

list nntp

Displays the properties of all of the NNTP monitors.

run nntp my_nntp destination 10.10.10.10:80

Runs a one-shot test of the custom monitor my_nntp against a target node at 10.10.10.10:80.

stop nntp my_nntp

Cancels a one-shot test of the custom monitor my_nntp in progress.

show nntp my_nntp test-result

Displays the result of the most recent one-shot test of the custom monitor my_nntp.
```

OPTIONS

app-service
Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug
Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from
Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is nntp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

newsgroup

Specifies the name of the newsgroup that you are monitoring, for example alt.car.mercedes. The default value is none.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor nntp(1)

ltm monitor none

NAME

none - Disables monitoring for a node or pool member.

MODULE

ltm monitor

SYNTAX

Apply the none monitor using the syntax in the following sections. The "none" attribute is indicated when viewing the listed node.

MODIFY

modify ltm node [[name] | [glob] | [regex]] ...] monitor none

DISPLAY

list ltm node monitor

DESCRIPTION

To disable monitoring on the node or pool member, set the monitor to none.

EXAMPLES

modify ltm pool [name] monitor none

Apply a none monitor on a pool.

```
modify ltm pool members modify { [name] { monitor none } }
```

Apply a none monitor on a pool member.

```
modify ltm node [name] monitor none
```

Apply a none monitor on a node.

```
list ltm node monitor
```

Displays the monitor applied to all nodes.

```
list ltm node [name] monitor
```

Displays the monitor applied to a specific node.

```
list ltm pool p1 members { all { monitor } }
```

Displays the monitor applied to all pools' members.

```
list ltm pool [name] members { monitor }
```

Displays the monitor applied to a pool's members.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2017-01-20 ltm monitor none(1)

ltm monitor oracle

NAME

oracle - Configures an Oracle(r) monitor.

MODULE

ltm monitor

SYNTAX

Configure the oracle component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create oracle [name]

modify oracle [name]

options:

app-service [[string] | none]

count [integer]

database [[name] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

manual-resume [enabled | disabled]

password [none | [password]]

recv [none | [string]]

recv-column [none | [string]]

recv-row [none | [string]]

send [none | [string]]

time-until-up [integer]

timeout [integer]

up-interval [integer]

username [[name] | none]

edit oracle [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list oracle

list oracle [[[name] | [glob] | [regex]] ...]

show oracle [[[name] | [glob] | [regex]] ...]

show running-config oracle

show running-config oracle [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete oracle [name]

Note: You cannot delete default monitors.

RUN

run oracle [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop oracle [name]

DESCRIPTION

You can use the oracle component to configure a custom monitor, or you can use the default Oracle monitor that the Local Traffic Manager provides. This type of monitor verifies Oracle database services.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create oracle my_oracle defaults-from oracle
```

Creates a monitor named my_oracle that inherits properties from the default Oracle monitor.

```
list oracle
```

Displays the properties of all of the Oracle monitors.

```
run oracle my_oracle destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_oracle against a target node at 10.10.10.10:80.

stop oracle my_oracle

Cancels a one-shot test of the custom monitor my_oracle in progress.

show oracle my_oracle test-result

Displays the result of the most recent one-shot test of the custom monitor my_oracle.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

count

Specifies the number of monitor probes after which the connection to the database will be terminated. Count value of zero indicates that the connection will never be terminated. The default value is zero.

database

Specifies the name of the database with which the monitor attempts to communicate. The proper format for database name is ::. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is oracle.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 30 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

`recv` Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

`recv-column`

Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

`recv-row`

Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`send` Specifies the SQL query that the monitor sends to the target database, for example, `SELECT count(*) FROM mytable`.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

`test-result`

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

`time-until-up`

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

`up-interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

`username`

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `run`, `show`, `stop`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor oracle(1)

ltm monitor pop3

NAME

`pop3` - Configures a Post Office Protocol (POP3) monitor.

MODULE

ltm monitor

SYNTAX

Configure the `pop3` component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY
create pop3 [name]
modify pop3 [name]
options:
 app-service [[string] | none]
 debug [no | yes]
 defaults-from [name]
 description [string]
 destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
 interval [integer]
 manual-resume [enabled | disabled]
 password [none | [password]]
 time-until-up [integer]
 timeout [integer]
 up-interval [integer]
 username [[name] | none]

edit pop3 [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties

DISPLAY
list pop3
list pop3 [[[name] | [glob] | [regex]] ...]
show pop3 [[[name] | [glob] | [regex]] ...]
show running-config pop3
show running-config pop3 [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties
 one-line
 partition
 test-result

DELETE
delete pop3 [name]

Note: You cannot delete default monitors.

RUN
run pop3 [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP
stop pop3 [name]

DESCRIPTION

You can use the pop3 component to configure a custom monitor, or you can use the default POP3 monitor that the Local Traffic Manager provides. This type of monitor verifies the POP3 service by attempting to connect to a pool, pool member, or virtual server, log on as the specified user, and log off.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create pop3 my_pop3 defaults-from pop3
```

Creates a monitor named my_pop3 that inherits properties from the default POP3 monitor.

```
list pop3
```

Displays the properties of all of the POP3 monitors.

```
run pop3 my_pop3 destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_pop3 against a target node at 10.10.10.10:80.

```
stop pop3 my_pop3
```

Cancels a one-shot test of the custom monitor my_pop3 in progress.

```
show pop3 my_pop3 test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_pop3.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and

labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is pop3.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor pop3(1)

ltm monitor postgresql

NAME

postgresql - Configures a PostgreSQL(r) monitor.

MODULE

ltm monitor

SYNTAX

Configure the postgresql component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create postgresql [name]

modify postgresql [name]

options:

app-service [[string] | none]

count [integer]

database [[name] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

manual-resume [enabled | disabled]

password [none | [password]]

recv [none | [string]]

recv-column [none | [string]]

recv-row [none | [string]]

send [none | [string]]

time-until-up [integer]

timeout [integer]

up-interval [integer]

username [[name] | none]

edit postgresql [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list postgresql

list postgresql [[[name] | [glob] | [regex]] ...]

show postgresql [[[name] | [glob] | [regex]] ...]

show running-config postgresql

show running-config postgresql [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete postgresql [name]

Note: You cannot delete default monitors.

RUN

```
run postgresql [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

STOP

```
stop postgresql [name]
```

DESCRIPTION

You can use the postgresql component to configure a custom monitor, or you can use the default PostgreSQL monitor that the Local Traffic Manager provides. This type of monitor verifies PostgreSQL-based services.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create postgresql my_postgresql defaults-from postgresql
```

Creates a monitor named my_postgresql that inherits properties from the default PostgreSQL monitor.

```
list postgresql
```

Displays the properties of all of the PostgreSQL monitors.

```
run postgresql my_postgresql destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_postgresql against a target node at 10.10.10.10:80.

```
stop postgresql my_postgresql
```

Cancels a one-shot test of the custom monitor my_postgresql in progress.

```
show postgresql my_postgresql test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_postgresql.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

count

Specifies the number of monitor probes after which the connection to the database will be terminated. Count value of zero indicates that the connection will never be terminated. The default value is zero.

database

Specifies the name of the database with which the monitor attempts to communicate. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is postgresql.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and

port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval
Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 30 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume
Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition
Displays the administrative partition within which the component resides.

password
Specifies the password if the monitored target requires authentication. The default value is none.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in a field in your database. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-column
Specifies the column in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

recv-row
Specifies the row in the database where the system expects the specified Receive String to be located. Specify this option only if you configure the send and recv options. The default value is none.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

send Specifies the SQL query that the monitor sends to the target database, for example, SELECT count(*) FROM mytable.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

test-result
Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up
Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout
Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 91 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval
Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username
Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor postgresql(1)

ltm monitor radius-accounting

NAME

radius-accounting - Configures a RADIUS accounting monitor for the BIG-IP(r) Local Traffic Manager.

MODULE

ltm monitor

SYNTAX

Configure the radius-accounting component within the ltm monitor module using the syntax shown in the following sections.

CREATE/MODIFY

create radius-accounting [name]

modify radius-accounting [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [[name] | none]

description [string]

destination [ip address]

interval [integer]

manual-resume [disabled | enabled]

nas-ip-address [ip address]

secret [string]

time-until-up [integer]

timeout [integer]

up-interval [integer]

username [none | [string]]

edit radius-accounting [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list radius-accounting

list radius-accounting [[[name] | [glob] | [regex]] ...]

show radius-accounting [[[name] | [glob] | [regex]] ...]

show running-config radius-accounting

show running-config radius-accounting [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete radius-accounting [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the radius-accounting component to configure a custom monitor, or you can use the default RADIUS accounting monitor that the Local Traffic Manager provides. This type of monitor provides information about the usage of the RADIUS service for accounting purposes.

EXAMPLES

```
create radius-accounting my_radius_acct defaults-from radius_accounting
```

Creates a monitor named my_radius_acct that inherits properties from the default RADIUS accounting monitor.

```
list radius-accounting
```

Displays the properties of all of the RADIUS accounting monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is radius.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

nas-ip-address

Specifies the network access server IP address that the system uses to identify itself to the RADIUS server. Using this option, multiple BIG-IP(r) systems can appear as a single network access device to the RADIUS server. The default value is none.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Specifies the secret the monitor needs to communicate with the resource. The default value is none.

time-until-up

Specifies the amount of time in seconds after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, ltm pool, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2014. All rights reserved.

BIG-IP 2017-04-05 ltm monitor radius-accounting(1)

ltm monitor radius

NAME

radius - Configures a Remote Access Dial-in User Service (RADIUS) monitor.

MODULE

ltm monitor

SYNTAX

Configure the radius component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create radius [name]

modify radius [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

manual-resume [enabled | disabled]

nas-ip-address [[ip address] | none]

password [none | [password]]

secret [none | [secret]]

time-until-up [integer]

timeout [integer]

up-interval [integer]

username [[name] | none]

edit radius [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list radius

list radius [[[name] | [glob] | [regex]] ...]

show radius [[[name] | [glob] | [regex]] ...]

show running-config radius

show running-config radius [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete radius [name]

Note: You cannot delete default monitors.

RUN

```
run radius [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

STOP

```
stop radius [name]
```

DESCRIPTION

You can use the radius component to configure a custom monitor, or you can use the default RADIUS monitor that the Local Traffic Manager provides. This type of monitor verifies the RADIUS service by attempting to authenticate the specified user.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create radius my_radius defaults-from radius
```

Creates a monitor named my_radius that inherits properties from the default RADIUS monitor.

```
list radius
```

Displays the properties of all of the RADIUS monitors.

```
run radius my_radius destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_radius against a target node at 10.10.10.10:80.

```
stop radius my_radius
```

Cancels a one-shot test of the custom monitor my_radius in progress.

```
show radius my_radius test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_radius.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is radius.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression

syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

nas-ip-address

Specifies the network access server IP address that the system uses to identify itself to the RADIUS server. With this option, multiple BIG-IP systems can appear as a single network access device to the RADIUS server. The default value is none.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Specifies the secret the monitor must use when contacting the resource. The default value is none.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the username, if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

Itm monitor real-server

NAME

real-server - Configures a RealServer(r) monitor.

MODULE

itm monitor

SYNTAX

Configure the real-server component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

create real-server [name]

modify real-server [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

interval [integer]

metrics [[metrics] | none]

time-until-up [integer]

timeout [integer]

edit real-server [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list real-server

list real-server [[[name] | [glob] | [regex]] ...]

show real-server [[[name] | [glob] | [regex]] ...]

show running-config real-server

show running-config real-server [[[name] | [glob] | [regex]] ...]

options:

agent

all-properties

command

method

non-default-properties

one-line

partition

DELETE

delete real-server [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the real-server component to configure a custom monitor, or you can use the default RealServer monitor that the Local Traffic Manager provides. This type of monitor checks the performance of a pool, pool member, or virtual server that is running the RealServer data collection agent, and then dynamically load balances traffic accordingly.

EXAMPLES

```
create real-server my_real-server defaults-from real_server
```

Creates a monitor named my_real-server that inherits properties from the default RealServer monitor.

```
list real-server
```

Displays the properties of all of the RealServer monitors.

OPTIONS

agent

Displays the agent for the monitor. The default agent is Mozilla/4.0 (compatible: MSIE 5.0; Windows NT). You cannot modify the agent.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

command

Displays the command that the system uses to obtain the metrics from the resource. See the documentation for this resource for information on available commands. You cannot modify the command.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is real-server.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 5 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

method

Displays the GET method. You cannot modify the method.

metrics

Specifies the performance metrics that the commands collect from the target. The default value is ServerBandwidth:1.5, CPUPercentUsage, MemoryUsage, TotalClientCount.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2017-04-05 ltm monitor real-server(1)

ltm monitor rpc

NAME

rpc - Configures a Remote Procedure Call (RPC) monitor.

MODULE

ltm monitor

SYNTAX

Configure the rpc component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create rpc [name]

modify rpc [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

interval [integer]

manual-resume [enabled | disabled]

mode [tcp | udp]

program-number [[integer] | none]

time-until-up [integer]

timeout [integer]

up-interval [integer]

version-number [[integer] | none]

edit rpc [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

```
list rpc
list rpc [ [ [name] | [glob] | [regex] ] ... ]
show rpc [ [ [name] | [glob] | [regex] ] ... ]
show running-config rpc
show running-config rpc [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

DELETE

```
delete rpc [name]
```

Note: You cannot delete default monitors.

RUN

```
run rpc [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

STOP

```
stop rpc [name]
```

DESCRIPTION

You can use the rpc component to configure a custom monitor, or you can use the default RPC monitor that the Local Traffic Manager provides. This type of monitor queries the RPC server, and verifies the availability of a given program.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create rpc my_rpc defaults-from rpc
```

Creates a monitor named my_rpc that inherits properties from the default RPC monitor.

```
list rpc
```

Displays the properties of all of the RPC monitors.

```
run rpc my_rpc destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_rpc against a target node at 10.10.10.10:80.

```
stop rpc my_rpc
```

Cancels a one-shot test of the custom monitor my_rpc in progress.

```
show rpc my_rpc test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_rpc.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is rpc.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The

default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

mode Specifies the protocol that the monitor uses to communicate with the target. The default value is tcp.

The options are:

tcp Specifies that the monitor uses the TCP protocol to communicate with the target object.

udp Specifies that the monitor uses the UDP protocol to communicate with the target object.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

program-number

Specifies the number of the program for which you want the monitor to verify availability. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

version-number

Specifies the number of the version for which you want the monitor to verify availability. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor rpc(1)

ltm monitor sasp

NAME

sasp - Configures a Server Application State Protocol (SASP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the sasp component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create sasp [name]

modify sasp [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

interval [integer]

mode [pull | push]

primary-address [ip address]

protocol [tcp | udp]

secondary-address [[ip address] | none]

service [none | [port]]

time-until-up [integer]

timeout [integer]

edit sasp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list sasp

list sasp [[[name] | [glob] | [regex]] ...]

show sasp [[[name] | [glob] | [regex]] ...]

show running-config sasp

show running-config sasp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete sasp [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the sasp component to configure a custom monitor, or you can use the default FTP monitor that the Local Traffic Manager provides. This type of monitor verifies the availability of IBM Group Workload Managers network resources.

EXAMPLES

```
create sasp my_sasp defaults-from sasp
```

Creates a monitor named my_sasp that inherits properties from the default SASP monitor.

```
list sasp
```

Displays the properties of all of the SASP monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `sasp`.

description

User defined description.

glob Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

interval

Specifies the frequency at which the system issues the monitor check. The default value is `auto`.

mode Specifies whether the load balancer should send Get Weight Request messages (pull) or receive Send Weights messages (push) from the GWM. The default mode is `push`.

name Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

partition

Displays the administrative partition within which the component resides.

primary-address

Specifies the IP address of the primary Group Workload Manager.

protocol

Specifies the protocol that the monitor uses to communicate with the target. The default value is `tcp`.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

secondary-address

Specifies the IP address of the secondary Group Workload Manager.

service

Specifies the port through which the SASP monitor communicates with the Group Workload Manager. The default port is `3860`.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 100 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2017-11-17 ltm monitor sasp(1)

Itm monitor scripted

NAME

`scripted` - Configures a Scripted monitor.

MODULE

ltm monitor

SYNTAX

Configure the scripted component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY
create scripted [name]
modify scripted [name]
options:
 app-service [[string] | none]
 debug [no | yes]
 defaults-from [name]
 description [string]
 destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
 filename [[filename] | none]
 interval [integer]
 manual-resume [enabled | disabled]
 time-until-up [integer]
 timeout [integer]
 up-interval [integer]

edit scripted [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties

DISPLAY
list scripted
list scripted [[[name] | [glob] | [regex]] ...]
show scripted [[[name] | [glob] | [regex]] ...]
show running-config scripted
show running-config scripted [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties
 one-line
 partition
 test-result

DELETE
delete scripted [name]

Note: You cannot delete default monitors.

RUN
run scripted [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP
stop scripted [name]

DESCRIPTION

You can use the scripted component to configure a custom monitor, or you can use the default scripted monitor that the Local Traffic Manager provides.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create scripted my_scripted defaults-from scripted
```

Creates a monitor named my_scripted that inherits properties from the default scripted monitor.

```
list scripted
```

Displays the properties of all of the scripted monitors.

```
run scripted my_scripted destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_scripted against a target node at 10.10.10.10:80.

```
stop scripted my_scripted
```

Cancels a one-shot test of the custom monitor my_scripted in progress.

```
show scripted my_scripted test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_scripted.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is scripted.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

filename

Specifies the name of a file in the `/config/eav/` directory on the system. The user-created file contains the send and expect data that the monitor uses for the monitor check. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the `up-interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is

up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor scripted(1)

ltm monitor sip

NAME

sip - Configures a Session Initiation Protocol (SIP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the sip component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create sip [name]
modify sip [name]
options:
  app-service [[string] | none]
  cert [ [cert list] | none]
  cipherlist [string]
  compatibility [enabled | disabled]
  debug [ no | yes]
  defaults-from [name]
  description [string]
  destination [ [ ipv4 address[:port] ] | [ ipv6 address[.port] ] ]
  filter [any | none | status]
  filter-neg [any | none | status]
  headers [ [new line separated headers] | none]
  interval [integer]
  key [ [key] | none]
  manual-resume [enabled | disabled]
  mode [sips | tcp | tls | udp | mr-tls | mr-sips | mr-tcp | mr-udp | mr-sctp]
  request [none | [string] ]
  time-until-up [integer]
  up-interval [integer]
  username [ [name] | none]
```

```
edit sip [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list sip
list sip [ [ [name] | [glob] | [regex] ] ... ]
show sip [ [ [name] | [glob] | [regex] ] ... ]
show running-config sip
show running-config sip [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

DELETE

```
delete sip [name]
```

Note: You cannot delete default monitors.

RUN

```
run sip [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[.port] ] ] ]
```

STOP
stop sip [name]

DESCRIPTION

You can use the sip component to configure a custom monitor, or you can use the default SIP monitor that the Local Traffic Manager provides. This type of monitor checks the status of SIP Call-ID services on a device. The SIP protocol enables real-time messaging, voice, data, and video.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create sip my_sip defaults-from sip
```

Creates a monitor named my_sip that inherits properties from the default SIP monitor.

```
list sip
```

Displays the properties of all of the SIP monitors.

```
run sip my_sip destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_sip against a target node at 10.10.10.10:80.

```
stop sip my_sip
```

Cancels a one-shot test of the custom monitor my_sip in progress.

```
show sip my_sip test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_sip.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

cert Specifies a fully-qualified path for a client certificate that the monitor sends to the target SSL server. The default value is none. Currently ignored for "mr-" modes.

cipherlist

Specifies the list of ciphers for this monitor. The default value is DEFAULT:+SHA:+3DES:+kEDH. Currently ignored for "mr-" modes.

compatibility

Specifies, when enabled, that the SSL options setting (in OpenSSL) is set to ALL. The default value is enabled. Currently ignored for "mr-" modes.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks.

The default value is no. The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is sip.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

filter

Specifies the SIP status codes that the target can return to be considered up. By default the system always accepts status codes whose value is in the 100, 200 or 300s.

The options are:

any Specifies that the monitor accepts any SIP status codes.

none Specifies that the monitor does not accept any other SIP status codes. This is the default value.

status

Specifies one or more status codes that you want to add to the monitor.

filter-neg

Specifies the SIP status codes that the target can return to be considered down. By default the system always accepts status codes according to sip-monitor.filter. After checking that, the status code is checked against this key. If a code is also in sip-monitor.filter, the node is marked up.

The options are:

any Specifies that the monitor rejects all SIP status codes that are not in sip-monitor.filter.

none Specifies that the monitor does not specifically reject any other SIP status codes. This is the default value.

status

Specifies one or more status codes that you want to add to the monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

headers

Specifies the set of SIP headers in the SIP message that is sent to the target. Separate each header with a new line. The default value is none.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

key Specifies the key if the monitored target requires authentication. The default value is none. Currently ignored for "mr-" modes.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

mode Specifies the transport protocol that the monitor uses to communicate with the target. The default mode is udp. The options are:

sips Specifies that the monitor uses SIPS to communicate with the target.

tcp Specifies that the monitor uses TCP to communicate with the target.

tls Specifies that the monitor uses TLS to communicate with the target, and the SIP URI is SIPS.

udp Specifies that the monitor uses UDP to communicate with the target.

mr-tls

Specifies that the monitor uses TLS with the message-routing framework to communicate with the target. This is serverside SSL over TCP. Note: Customizing the SSL-related options (cert, cipherlist, compatibility, key) is currently ignored for this mode, and the monitor will operate with the default values.

mr-sips

Specifies that the monitor uses SIPS with the message-routing framework to communicate with the target. This is TLS mode with "sips" replacing "sip" in the SIP message headers. Note: Customizing the SSL-related options (cert, cipherlist, compatibility, key) is currently ignored for this mode, and the monitor will operate with the default values.

mr-tcp

Specifies that the monitor uses TCP with the message-routing framework to communicate with the target.

mr-udp

Specifies that the monitor uses UDP with the message-routing framework to communicate with the target.

`mr-sctp`

Specifies that the monitor uses SCTP with the message-routing framework to communicate with the target.

The modes beginning with "mr-" use an in-TMM monitor based on the message-routing framework (MRF). The other modes use the legacy monitor based on message-based load balancing (MLBL).

`name` Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`request`

Specifies the SIP request line in the SIP message that is sent to the target. The default value is none.

`test-result`

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

`time-until-up`

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is

`timeout`

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

`up-interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2016, 2017. All rights reserved.

BIG-IP 2019-10-28 ltm monitor sip(1)

ltm monitor smb

NAME

smb - Configures a Server Message Bloc (SMB)/Common Internet File System (CIFS) monitor.

MODULE

ltm monitor

SYNTAX

Configure the smb component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create smb [name]

modify smb [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]
destination [[ipv4 address[:port]] | [ipv6 address[:port]]]
get [none | [filename]]
interval [integer]
manual-resume [enabled | disabled]
password [none | [password]]
server [[NETBIOS name] | none]
service [[[name] | [integer]] | none]
time-until-up [integer]
timeout [integer]
up-interval [integer]
username [[name] | none]

edit smb [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list smb

list smb [[[name] | [glob] | [regex]] ...]

show smb [[[name] | [glob] | [regex]] ...]

show running-config smb

show running-config smb [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition
test-result

DELETE

delete smb [name]

Note: You cannot delete default monitors.

RUN

run smb [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP

stop smb [name]

DESCRIPTION

You can use the smb component to configure a custom monitor, or you can use the default SMB monitor that the Local Traffic Manager provides. This type of monitor verifies the availability of an SMB/CIFS server. You can use this type of monitor to either check the availability of the server as a whole, the availability of a specific service on the server, or the availability of a specific file used by a service.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

create smb my_smb defaults-from smb

Creates a monitor named my_smb that inherits properties from the default SMB monitor.

list smb

Displays the properties of all of the SMB monitors.

run smb my_smb destination 10.10.10.10:80

Runs a one-shot test of the custom monitor my_smb against a target node at 10.10.10.10:80.

stop smb my_smb

Cancels a one-shot test of the custom monitor my_smb in progress.

show smb my_smb test-result

Displays the result of the most recent one-shot test of the custom monitor my_smb.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

`no` Specifies that the system does not redirect error messages and additional information related to this monitor.

`yes` Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `smb`.

`description`

User defined description.

`destination`

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

`get` Specifies a file associated with a service. The default value is none.

The monitor uses the relative path to the service itself when attempting to locate the file. You are not required to specify a value for this option; however, if you elect to use this option you must also specify a value for the service option.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`interval`

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the `up-interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

`manual-resume`

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

`partition`

Displays the administrative partition within which the component resides.

`password`

Specifies the password if the monitored target requires authentication. The default value is none.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`server`

Specifies the NetBIOS name of the SMB/CIFS server for which this monitor checks for availability. You must specify a server for this monitor to function. The default value is none.

`service`

Specifies a specific service on the SMB/CIFS for which you want to verify availability. The default value is none.

`test-result`

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

`time-until-up`

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the

node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the user name if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2016, 2017. All rights reserved.

BIG-IP 2017-08-16 Itm monitor smb(1)

Itm monitor smtp

NAME

smtp - Configures a Simple Mail Transport Protocol (SMTP) monitor.

MODULE

itm monitor

SYNTAX

Configure the smtp component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

create smtp [name]

modify smtp [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

domain [[name] | none]

interval [integer]

manual-resume [enabled | disabled]

time-until-up [integer]

timeout [integer]

up-interval [integer]

edit smtp [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list smtp

list smtp [[name] | [glob] | [regex]] ...]

show smtp [[name] | [glob] | [regex]] ...]

show running-config smtp

show running-config smtp [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE
delete smtp [name]

Note: You cannot delete default monitors.

RUN
run smtp [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP
stop smtp [name]

DESCRIPTION

You can use the smtp component to configure a custom monitor, or you can use the default SMTP monitor that the Local Traffic Manager provides. This type of monitor checks the status of a pool, pool member, or virtual server by issuing standard SMTP commands.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create smtp my_smtp defaults-from smtp
```

Creates a monitor named my_smtp that inherits properties from the default SMTP monitor.

```
list smtp
```

Displays the properties of all of the SMTP monitors.

```
run smtp my_smtp destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_smtp against a target node at 10.10.10.10:80.

```
stop smtp my_smtp
```

Cancels a one-shot test of the custom monitor my_smtp in progress.

```
show smtp my_smtp test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_smtp.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is smtp.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the

destination option for the specified custom monitor.

domain

Specifies the domain name to check, for example, bigipinternal.com. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2016-2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor smtp(1)

ltm monitor snmp-dca-base

NAME

snmp-dca-base - Configures a base Simple Network Management Protocol (SNMP) Data Center Audit monitor.

MODULE

ltm monitor

SYNTAX

Configure the snmp-dca-base component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create snmp-dca-base [name]
modify snmp-dca-base [name]
options:
  app-service [[string] | none]
  community [ [name] | none]
  cpu-coefficient [ [integer] | none]
  defaults-from [name]
  description [string]
  interval [integer]
  time-until-up [integer]
  timeout [integer]
  user-defined [ [name] [value] | [name] none ]
  version [ [integer] | none]

edit snmp-dca-base [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list snmp-dca-base
list snmp-dca-base [ [ [name] | [glob] | [regex] ] ... ]
show snmp-dca-base [ [ [name] | [glob] | [regex] ] ... ]
show running-config snmp-dca-base
show running-config snmp-dca-base [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete snmp-dca-base [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the snmp-dca-base component to configure a custom monitor, or you can use the default base SNMP DCA monitor that the Local Traffic Manager provides. This type of monitor checks the performance of a server running an SNMP agent such as UC Davis. Use this monitor only when you want the load balancing destination to be based solely on user data, and not CPU, memory or disk use.

EXAMPLES

```
create snmp-dca-base my_snmp-dca-base defaults-from snmp_dca_base
```

Creates a monitor named my_snmp-dca-base that inherits properties from the default base SNMP DCA monitor.

```
list snmp-dca-base
```

Displays the properties of all of the base SNMP DCA monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

community

Specifies the community name that the BIG-IP system must use to authenticate with the host server through SNMP. The default value is public.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is snmp_dca_base.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 10 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 30 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

user-defined

Specifies any user-defined command-line arguments and variables that the external program requires. Use the following syntax to specify a user defined parameter.

```
modify external my_external user-defined my_param_name my_param_value
```

Use the following syntax to remove a user defined parameter.

```
modify external my_external user-defined my_param_name none
```

version

Specifies the version of SNMP that the host server uses. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2017-04-05 Itm monitor snmp-dca-base(1)

Itm monitor snmp-dca

NAME

snmp-dca - Configures a Simple Network Management Protocol (SNMP) Data Center Audit monitor.

MODULE

itm monitor

SYNTAX

Configure the snmp component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

```
create snmp-dca [name]
```

```
modify snmp-dca [name]
```

options:

```
agent-type [generic | other | win2000 | ucd]
```

```
app-service [[string] | none]
```

```
community [ [name] | none]
```

```
cpu-coefficient [ [integer] | none]
```

```
cpu-threshold [none | [integer] ]
```

```
defaults-from [name]
```

```
description [string]
```

```
disk-coefficient [ [integer] | none]
```

```
disk-threshold [none | [integer] ]
```

```
interval [integer]
```

```
memory-coefficient [ [integer] | none]
```

```
memory-threshold [none | [integer] ]
```

```
time-until-up [integer]
```

```
timeout [integer]
```

```
user-defined
```

```
version [ [integer] | none]
```

```
edit snmp-dca [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list snmp-dca
list snmp-dca [ [ [name] | [glob] | [regex] ] ... ]
show snmp-dca [ [ [name] | [glob] | [regex] ] ... ]
show running-config snmp-dca
show running-config snmp-dca [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete snmp-dca [name]
```

Note: You cannot delete default monitors.

DESCRIPTION

You can use the `snmp-dca` component to configure a custom monitor, or you can use the default SNMP DCA monitor that the Local Traffic Manager provides. This type of monitor checks the performance of a server running an SNMP agent such as UC Davis, for the purpose of load balancing traffic to that server.

EXAMPLES

```
create snmp-dca my_snmp-dca defaults-from snmp_dca
```

Creates a monitor named `my_snmp-dca` that inherits properties from the default SNMP DCA monitor.

```
list snmp-dca
```

Displays the properties of all of the SNMP DCA monitors.

OPTIONS

`agent-type`

Specifies the type of agent. The default value is `ucd`.

`app-service`

Specifies the name of the application service to which the monitor belongs. The default value is `none`.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

`community`

Specifies the community name that the BIG-IP system must use to authenticate with the host server through SNMP. The default value is `public`.

`cpu-coefficient`

Specifies the coefficient that the system uses to calculate the weight of the CPU threshold in the dynamic ratio load balancing algorithm. The default value is `1.5`.

`cpu-threshold`

Specifies the maximum acceptable CPU usage on the target server. The default value is `80` percent.

`defaults-from`

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `snmp_dca`.

`description`

User defined description.

`disk-coefficient`

Specifies the coefficient that the system uses to calculate the weight of the disk threshold in the dynamic ratio load balancing algorithm. The default value is `2.0`.

`disk-threshold`

Specifies the maximum acceptable disk usage on the target server. The default value is `90` percent.

`glob` Displays the items that match the `glob` expression. See `help glob` for a description of `glob` expression syntax.

`interval`

Specifies the frequency at which the system issues the monitor check. The default value is `10` seconds.

`memory-coefficient`

Specifies the coefficient that the system uses to calculate the weight of the memory threshold in the dynamic ratio load balancing algorithm. The default value is `1.0`.

`memory-threshold`

Specifies the maximum acceptable memory usage on the target server. The default value is `70` percent.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an `@` sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See `help regex` for

a description of regular expression syntax.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 30 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

user-defined

Specifies attributes for a monitor that you define. The default value is none.

version

Specifies the version of SNMP that the host server uses. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2017-04-05 ltm monitor snmp-dca(1)

ltm monitor soap

NAME

soap - Configures a Simple Object Access Protocol (SOAP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the soap component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create soap [name]

modify soap [name]

options:

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

expect-fault [no | yes]

interval [integer]

manual-resume [enabled | disabled]

method [string]

namespace [[name] | none]

parameter-name [[name] | none]

parameter-type [bool | int | long | string]

parameter-value [none | [integer] | [string]]

password [none | [password]]

protocol [http | https]

return-type [bool | char | double | int | long | short | string]

return-value [none | [integer] | [string]]

soap-action [string]

time-until-up [integer]

timeout [integer]

up-interval [integer]

url-path [none | [string]]

username [[name] | none]

edit soap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list soap

```
list soap [ [ [name] | [glob] | [regex] ] ... ]
show soap [ [ [name] | [glob] | [regex] ] ... ]
show running-config soap
show running-config soap [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

```
DELETE
delete soap [name]
```

Note: You cannot delete default monitors.

```
RUN
run soap [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

```
STOP
stop soap [name]
```

DESCRIPTION

You can use the soap component to configure a custom monitor, or you can use the default SOAP monitor that the Local Traffic Manager provides. This type of monitor tests a Web service based on SOAP.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create soap my_soap defaults-from soap
```

Creates a soap monitor that inherits values from the system default SOAP monitor.

```
list soap
```

Displays the properties of all of the SOAP monitors.

```
run soap my_soap destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_soap against a target node at 10.10.10.10:80.

```
stop soap my_soap
```

Cancels a one-shot test of the custom monitor my_soap in progress.

```
show soap my_soap test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_soap.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.

defaults-from

Specifies the type of monitor you want to use to create the new monitor. The default value is soap.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member

and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

expect-fault

Specifies whether the value of the method option causes the monitor to expect a SOAP fault message. The default value is no.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

method

Specifies the method by which the monitor contacts the resource.

name

Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

namespace

Specifies the name space for the Web service you are monitoring, for example, http://example.com/. The default value is none.

parameter-name

If the method has a parameter, specifies the name of that parameter. The default value is none.

parameter-type

Specifies the parameter type. The default value is bool.

parameter-value

Specifies the value for the parameter. The default value is none.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

protocol

Specifies the protocol that the monitor uses to communicate with the target, http or https. The default value is http.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

return-type

Specifies the type for the returned parameter. The default value is bool.

return-value

Specifies the value for the returned parameter. The default value is none.

soap-action

Specifies the value for the SOAPAction header. The default value is the empty string.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds

with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

url-path

Specifies the URL for the Web service that you are monitoring, for example, /services/myService.aspx. The default value is none.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

username

Specifies the user name if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor soap(1)

ltm monitor tcp-echo

NAME

tcp-echo - Configures a Transmission Control Protocol (TCP) Echo monitor.

MODULE

ltm monitor

SYNTAX

Configure the tcp-echo component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create tcp-echo [name]

modify tcp-echo [name]

options:

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

app-service [[string] | none]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

manual-resume [enabled | disabled]

time-until-up [integer]

timeout [integer]

transparent [disabled | enabled]

up-interval [integer]

edit tcp-echo [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tcp-echo

list tcp-echo [[[name] | [glob] | [regex]] ...]

show tcp-echo [[[name] | [glob] | [regex]] ...]

show running-config tcp-echo

show running-config tcp-echo [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE
delete tcp-echo [name]

Note: You cannot delete default monitors.

RUN
run tcp-echo [name] [destination [[ipv4 address[:port]] | [ipv6 address[:port]]]]

STOP
stop tcp-echo [name]

DESCRIPTION

You can use the tcp-echo component to configure a custom monitor, or you can use the default TCP Echo monitor that the Local Traffic Manager provides. This type of monitor checks the status of a resource, using TCP Echo.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create tcp-echo my_tcp-echo defaults-from tcp_echo
```

Creates a monitor named my_tcp-echo that inherits properties from the default TCP Echo monitor.

```
list tcp-echo
```

Displays the properties of all of the TCP Echo monitors.

```
run tcp-echo my_tcp-echo destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_tcp-echo against a target node at 10.10.10.10:80.

```
stop tcp-echo my_tcp-echo
```

Cancels a one-shot test of the custom monitor my_tcp-echo in progress.

```
show tcp-echo my_tcp-echo test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_tcp-echo.

OPTIONS

adaptive

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

adaptive-divergence-type

Specifies whether the adaptive-divergence-value is relative or absolute.

adaptive-divergence-value

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%). If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool member or node would be marked down.)

adaptive-limit

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

adaptive-sampling-timespan

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is tcp_echo.

description

User defined description.

destination

Specifies the IP address of the resource that is the destination of this monitor. The default value is *.

Possible values are:

* Specifies to perform a health check on the IP address of the node.

IP address

Specifies to perform a health check on the IP address that you specify, and mark the associated node up or down accordingly.

IP address (with the transparent option enabled)

Specifies to perform a health check on the IP address that you specify, route the check through the IP address of the associated node, and mark the IP address of the associated node up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016-2017. All rights reserved.

Itm monitor tcp-half-open

NAME

tcp-half-open - Configures a Transmission Control Protocol (TCP) Half Open monitor.

MODULE

itm monitor

SYNTAX

Configure the tcp-half-open component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

create tcp-half-open [name]

modify tcp-half-open [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

interval [integer]

manual-resume [enabled | disabled]

time-until-up [integer]

timeout [integer]

transparent [disabled | enabled]

up-interval [integer]

edit tcp-half-open [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tcp-half-open

list tcp-half-open [[[name] | [glob] | [regex]] ...]

show tcp-half-open [[[name] | [glob] | [regex]] ...]

show running-config tcp-half-open

show running-config tcp-half-open [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

test-result

DELETE

delete tcp-half-open [name]

Note: You cannot delete default monitors.

RUN

run tcp-half-open [name] [destination [[ipv4 address[:port]] | [ipv6 address[.port]]]]

STOP

stop tcp-half-open [name]

DESCRIPTION

You can use the tcp-half-open component to configure a custom monitor, or you can use the default TCP Half Open monitor that the Local Traffic Manager provides.

For more information about configuring monitors, refer to the Configuration Guide for BIG-IP(r) Local Traffic Manager(r).

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create tcp-half-open my_tcp-half-open defaults-from tcp_half_open
```

Creates a monitor named my_tcp-half-open that inherits properties from the default TCP Half Open monitor.

```
list tcp-half-open
```

Displays the properties of all of the TCP Half Open monitors.

run tcp-half-open my_tcp-half-open destination 10.10.10.10:80

Runs a one-shot test of the custom monitor my_tcp-half-open against a target node at 10.10.10.10:80.

stop tcp-half-open my_tcp-half-open

Cancels a one-shot test of the custom monitor my_tcp-half-open in progress.

show tcp-half-open my_tcp-half-open test-result

Displays the result of the most recent one-shot test of the custom monitor my_tcp-half-open.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is tcp_half_open.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port you specify, route the check through the IP address and port supplied by the pool member, and mark the pool member (the gateway) up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012, 2017. All rights reserved.

BIG-IP 2017-08-16 Itm monitor tcp-half-open(1)

Itm monitor tcp

NAME

tcp - Configures a Transmission Control Protocol (TCP) monitor.

MODULE

itm monitor

SYNTAX

Configure the tcp component within the Itm monitor module using the syntax in the following sections.

CREATE/MODIFY

create tcp [name]

modify tcp [name]

options:

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

app-service [[string] | none]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[.port]]]

interval [integer]

ip-dscp [integer]

manual-resume [enabled | disabled]

recv [none | [string]]

recv-disable [none | [string]]

reverse [enabled | disabled]

send [none | [string]]

time-until-up [integer]

timeout [integer]

transparent [disabled | enabled]

up-interval [integer]

edit tcp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tcp

list tcp [[[name] | [glob] | [regex]] ...]

show tcp [[[name] | [glob] | [regex]] ...]

```
show running-config tcp
show running-config tcp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

```
DELETE
delete tcp [name]
```

Note: You cannot delete default monitors.

```
RUN
run tcp [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

```
STOP
stop tcp [name]
```

DESCRIPTION

You can use the tcp component to configure a custom monitor, or you can use the default TCP monitor that the Local Traffic Manager provides.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create tcp my_tcp defaults-from tcp
```

Creates a monitor named my_tcp that inherits properties from the default TCP monitor.

```
list tcp
```

Displays the properties of all of the TCP monitors.

```
run tcp my_tcp destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_tcp against a target node at 10.10.10.10:80.

```
stop tcp my_tcp
```

Cancels a one-shot test of the custom monitor my_tcp in progress.

```
show tcp my_tcp test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_tcp.

OPTIONS

adaptive

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

adaptive-divergence-type

Specifies whether the adaptive-divergence-value is relative or absolute.

adaptive-divergence-value

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%). If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool member or node would be marked down.)

adaptive-limit

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

adaptive-sampling-timespan

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The

default value is tcp.

description
User defined description.

destination
Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is *.*.

Possible values are:

. Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port
Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port
Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)
Specifies to perform a health check on the server at the IP address and port you specify, route the check through the IP address and port supplied by the pool member, and mark the pool member (the gateway) up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval
Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

ip-dscp
Specifies the differentiated services code point (DSCP). DSCP is a 6-bit value in the Differentiated Services (DS) field of the IP header. It can be used to specify the quality of service desired for the packet. The valid range for this value is 0 to 63 (hex 0x0 to 0x3f). The default value is zero.

manual-resume
Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition
Displays the administrative partition within which the component resides.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names. If you do not specify a value for both the send and recv options, the monitor performs a simple service check and connect only.

recv-disable
Specifies a text string that the monitor looks for in the returned resource. If the text string is matched in the returned resource, the corresponding node or pool member is marked session disabled. The default value is none.

You specify a recv-disable string in the same way that you specify a recv string.

If you specify a recv-disable string, you must also specify a recv string. You cannot specify a recv-disable string, if the reverse option is enabled.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reverse
Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use this mode only if you configure both the send and recv options.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

send Specifies the text string that the monitor sends to the target object. The default setting is GET /, which retrieves a default HTML file for a web site.

To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html. Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if not null.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2016-2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor tcp(1)

ltm monitor udp

NAME

udp - Configures a User Datagram Protocol (UDP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the udp component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create udp [name]

modify udp [name]

options:

adaptive [enabled | disabled]

adaptive-divergence-type [relative | absolute]

adaptive-divergence-value [integer]

adaptive-limit [integer]

adaptive-sampling-timespan [integer]

app-service [[string] | none]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

interval [integer]

manual-resume [enabled | disabled]

```
recv [none | [string] ]
recv-disable [none | [string] ]
reverse [enabled | disabled]
send [none | [string] ]
time-until-up [integer]
timeout [integer]
transparent [disabled | enabled]
up-interval [integer]
```

```
edit udp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list udp
list udp [ [ [name] | [glob] | [regex] ] ... ]
show udp [ [ [name] | [glob] | [regex] ] ... ]
show running-config udp
show running-config udp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  test-result
```

DELETE

```
delete udp [name]
```

Note: You cannot delete default monitors.

RUN

```
run udp [name] [ destination [ [ ipv4 address[:port] ] | [ ipv6 address[:port] ] ] ]
```

STOP

```
stop udp [name]
```

DESCRIPTION

You can use the udp component to configure a custom monitor, or you can use the default UDP monitor that the Local Traffic Manager provides. This type of monitor verifies the UDP service by attempting to send UDP packets to a pool, pool member, or virtual server and receiving a reply.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an ltm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create udp my_udp defaults-from udp
```

Creates a monitor named my_udp that inherits properties from the default UDP monitor.

```
list udp
```

Displays the properties of all of the UDP monitors.

```
run udp my_udp destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_udp against a target node at 10.10.10.10:80.

```
stop udp my_udp
```

Cancels a one-shot test of the custom monitor my_udp in progress.

```
show udp my_udp test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_udp.

OPTIONS

adaptive

Specifies whether the adaptive feature is enabled for this monitor. Not all monitors support the adaptive feature.

adaptive-divergence-type

Specifies whether the adaptive-divergence-value is relative or absolute.

adaptive-divergence-value

Specifies how far from mean latency each monitor probe is allowed to be. If adaptive-divergence-type is relative, this value is a percentage deviation from mean (e.g. 50 would indicate the probe is allowed to exceed the mean latency by 50%.) If adaptive-divergence-type is absolute, this value is an offset from mean in milliseconds (e.g. 250 would indicate the probe is allowed to exceed the mean latency by 250 ms.) A probe that exceeds latency is counted the same as a probe that is not received, so in the typical scenario, it will require three missed latencies in a row to mark a pool member or node down (i.e. a 15-second interval with a 46-second timeout, would require three missed probes before the pool

member or node would be marked down.)

adaptive-limit

Specifies the hard limit, in milliseconds, which the probe is not allowed to exceed, regardless of the divergence value. For example, if this value is 500, then the probe latency may not exceed 500 ms even if that would still fall within the divergence value.

adaptive-sampling-timespan

Specifies the size of the sliding window, in seconds, which records probe history. For example, if this value is 300, then a sliding window of the last five minutes' probe history will be used for calculating probe mean latency and standard deviation.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is `udp`.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. Possible values are:

: Specifies to perform a health check on the IP address and port supplied by a pool member.

*:port

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

IP address:port

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

IP address:port (with the transparent option enabled)

Specifies to perform a health check on the server at the IP address and port you specify, route the check through the IP address and port supplied by the pool member, and mark the pool member (the gateway) up or down accordingly.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the `up-interval` option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

recv Specifies the text string that the monitor looks for in the returned resource. The default value is none.

recv-disable

Specifies a text string that the monitor looks for in the returned resource. If the text string is

matched in the returned resource, the corresponding node or pool member is marked session disabled. The default value is none.

The recv-disable string may be specified the same way a recv string may be specified.

If the recv-disable string is configured, the recv string must be non-empty. The recv-disable string may not be configured if reverse mode is enabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reverse

Specifies whether the monitor operates in reverse mode. When the monitor is in reverse mode, a successful check marks the monitored object down instead of up. You can use the this mode only if you configure both the send and recv options.

The default value is disabled, which specifies that the monitor does not operate in reverse mode. The enabled value specifies that the monitor operates in reverse mode.

send Specifies the text string that the monitor sends to the target object. The default value is GET /, which retrieves a default HTML file for a web site.

To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html. Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if it is not null. The default value is none.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

transparent

Specifies whether the monitor operates in transparent mode. Monitors in transparent mode can monitor pool members through firewalls. The default value is disabled.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2016-2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor udp(1)

ltm monitor virtual-location

NAME

virtual-location - Configures a Virtual Location monitor.

MODULE

ltm monitor

SYNTAX

Configure the virtual-location component within the ltm monitor module using the syntax shown in the following sections.

CREATE/MODIFY

create virtual-location [name]
modify virtual-location [name]

options:

app-service [[string] | none]
debug [no | yes]
defaults-from [name]
description [string]
interval [integer]
pool [name]
time-until-up [integer]
timeout [integer]
up-interval [integer]

edit virtual-location [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list virtual-location

list virtual-location [[[name] | [glob] | [regex]] ...]

show virtual-location [[[name] | [glob] | [regex]] ...]

show running-config virtual-location

show running-config virtual-location

[[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

DELETE

delete virtual-location [name]

Note: You cannot delete default monitors.

DESCRIPTION

The Virtual Location monitor will determine if a pool member which has a virtual IP is currently a local pool member with its arp entry existing on a local VLAN, or, a remote pool member with its ARP entry existing on a tunnel VLAN. If the pool member is local it will set the pool member's priority to 2. If the pool member is remote it will set the priority to 1 (a lower priority). The Virtual Location will always return up as the availability for the pool member. It is necessary to use an additional monitor to check the availability status of the pool member.

You can use the virtual-location component to configure a custom monitor, or you can use the default Virtual Location monitor that the Local Traffic Manager provides.

EXAMPLES

```
create virtual-location my_virtual-location defaults-from virtual_location pool aPool
```

Creates a monitor named my_virtual-location that inherits properties from the default Virtual Location monitor.

```
list virtual-location
```

Displays the properties of all of the Virtual Location monitors.

OPTIONS

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is no. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The options are no (specifies that the system does not redirect error messages and additional information related to this monitor.) and yes (specifies that the system redirects error messages and additional information to the /var/log/monitors/--.log file.)

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is virtual_location.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

pool Specifies the pool for the target pool member. This is a required argument.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, ltm pool, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010, 2012-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm monitor virtual-location(1)

ltm monitor wap

NAME

wap - Configures a Wireless Application Protocol (WAP) monitor.

MODULE

ltm monitor

SYNTAX

Configure the wap component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create wap [name]

modify wap [name]

options:

accounting-node [none | [RADIUS server name]]

accounting-port [[integer] | none]

app-service [[string] | none]

call-id [none | [RADIUS server 11 digit phone number]]

debug [no | yes]

defaults-from [name]

description [string]

destination [[ipv4 address[:port]] | [ipv6 address[:port]]]

framed-address [none | [RADIUS framed IP address]]
interval [integer]
manual-resume [enabled | disabled]
recv [none | [string]]
secret [none | [password]]
send [none | [string]]
server-id [none | [RADIUS NAS-ID]]
session-id [none | [RADIUS session ID]]
time-until-up [integer]
timeout [integer]
up-interval [integer]

edit wap [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list wap

list wap [[[name] | [glob] | [regex]] ...]

show wap [[[name] | [glob] | [regex]] ...]

show running-config wap

show running-config wap [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition
test-result

DELETE

delete wap [name]

Note: You cannot delete default monitors.

RUN

run wap [name] [destination [[ipv4 address[:port]] | [ipv6 address[.port]]]]

STOP

stop wap [name]

DESCRIPTION

You can use the wap component to configure a custom monitor, or you can use the default WAP monitor that the Local Traffic Manager provides. This type of monitor requests the URL specified in the send option, and finds the string specified in the recv option somewhere in the data returned by the URL response.

You can test a custom monitor configuration against a specified target destination by using the run command, and view the results of such a test by using the show command with the test-result option.

The following user roles (in addition to the root user) have permissions to run and stop an Itm monitor test:

admin, application-editor, manager, operator, resource-admin

EXAMPLES

```
create wap my_wap defaults-from wap
```

Creates a monitor named my_wap that inherits properties from the default WAP monitor.

```
list wap
```

Displays the properties of all of the WAP monitors.

```
run wap my_wap destination 10.10.10.10:80
```

Runs a one-shot test of the custom monitor my_wap against a target node at 10.10.10.10:80.

```
stop wap my_wap
```

Cancels a one-shot test of the custom monitor my_wap in progress.

```
show wap my_wap test-result
```

Displays the result of the most recent one-shot test of the custom monitor my_wap.

OPTIONS

accounting-node

Specifies the RADIUS server that provides authentication for the WAP target. Note that if you configure the accounting-port option, but you do not configure the this option, the system assumes that the RADIUS server and the WAP server are the same system.

accounting-port

Specifies the port that the monitor uses for RADIUS accounting. The default value is none. A value of 0 (zero) disables RADIUS accounting.

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot

modify or delete the monitor. Only the application service can modify or delete the monitor.

call-id

Specifies the 11-digit phone number for the RADIUS server. The default value is none.

debug

Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. You can use the log information to help diagnose and troubleshoot unsuccessful health checks. The default value is no.

The options are:

no Specifies that the system does not redirect error messages and additional information related to this monitor.

yes Specifies that the system redirects error messages and additional information to the `/var/log/monitors/--.log` file.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is wap.

description

User defined description.

destination

Specifies the IP address and service port of the resource that is the destination of this monitor. The default value is `*:*`.

Possible values are:

`*:*` Specifies to perform a health check on the IP address and port supplied by a pool member.

`*:port`

Specifies to perform a health check on the server with the IP address supplied by the pool member and the port you specify.

`IP address:port`

Specifies to mark a pool member up or down based on the response of the server at the IP address and port you specify.

This option is required for the command run, unless an IP address and service port are specified in the destination option for the specified custom monitor.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

framed-address

Specifies the RADIUS framed IP address. The default value is none.

interval

Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.

Important: F5 Networks recommends that when you configure this option and the up-interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

manual-resume

Specifies whether the system automatically changes the status of a resource to up at the next successful monitor check. The default value of the manual-resume option is disabled.

Note that if you set the manual-resume option to enabled, you must manually mark the resource as up before the system can use it for load balancing connections.

name Specifies a unique name for the component. This option is required for the commands create, delete, modify, run and stop.

partition

Displays the administrative partition within which the component resides.

recv Specifies the text string that the monitor looks for in the returned resource. The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names. If you do not specify both a value for both the send and recv options, the monitor performs a simple service check and connect only. The default value is none.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

secret

Specifies the password the monitor needs to access the resource. The default value is none.

send Specifies the text string that the monitor sends to the target object. The default setting is GET /, which retrieves a default HTML file for a web site.

To retrieve a specific page from a web site, specify a fully-qualified path name, for example, GET /www/company/index.html. Since the string may have special characters, the system may require that the string be enclosed with single quotation marks.

If this value is null, then a valid connection suffices to determine that the service is up. In this case, the system does not need the recv option and ignores the option even if it is not null. The default value is none.

server-id

Specifies the RADIUS NAS-ID for this system when configuring a RADIUS server. The default value is none.

session-id

Specifies the RADIUS session identification number when configuring a RADIUS server. The default value is none.

test-result

Displays the result of the most recent one-shot test of the specified monitor(s), if any such test has been performed since BIG-IP was started.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 31 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

up-interval

Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Important: F5 Networks recommends that when you configure this option and the interval option, whichever value is greater be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.

SEE ALSO

create, delete, edit, glob, list, modify, regex, run, show, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-08-16 ltm monitor wap(1)

ltm monitor wmi

NAME

wmi - Configures a Windows Management Infrastructure (WMI) monitor.

MODULE

ltm monitor

SYNTAX

Configure the wmi component within the ltm monitor module using the syntax in the following sections.

CREATE/MODIFY

create wmi [name]

modify wmi [name]

options:

agent [string]

app-service [[string] | none]

command [[command] | none]

defaults-from [name]

description [string]

interval [integer]

metrics [[value] | none]

password [none | [password]]

time-until-up [integer]

timeout [integer]

url [none | [URL]]

username [[name] | none]

edit wmi [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list wmi

list wmi [[[name] | [glob] | [regex]] ...]

show wmi [[[name] | [glob] | [regex]] ...]

show running-config wmi

show running-config wmi [[[name] | [glob] | [regex]] ...]

options:

agent

all-properties

method

non-default-properties

one-line

partition

post

DELETE

delete wmi [name]

Note: You cannot delete default monitors.

DESCRIPTION

You can use the wmi component to configure a custom monitor, or you can use the default WMI monitor that the Local Traffic Manager provides. This type of monitor checks the performance of a pool, pool member, or virtual server that is running the WMI data collection agent, and then dynamically load balances traffic accordingly.

EXAMPLES

```
create wmi my_wmi defaults-from wmi
```

Creates a monitor named my_wmi that inherits properties from the default WMI monitor.

```
list wmi
```

Displays the properties of all of the WMI monitors.

OPTIONS

agent

Displays the agent for the monitor. The default agent is Mozilla/4.0 (compatible: MSIE 5.0; Windows NT).

app-service

Specifies the name of the application service to which the monitor belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the monitor. Only the application service can modify or delete the monitor.

command

Specifies the command that the system uses to obtain the metrics from the resource. See the documentation for this resource for information on available commands.

defaults-from

Specifies the name of the monitor from which you want your custom monitor to inherit settings. The default value is wmi.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interval

Specifies the frequency at which the system issues the monitor check. The default value is 5 seconds.

method

Displays the GET method. You cannot modify the method.

metrics

Specifies the performance metrics that the commands collect from the target. The default value is LoadPercentage, DiskUsage, PhysicalMemoryUsage:1.5, VirtualMemoryUsage:2.0.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

password

Specifies the password if the monitored target requires authentication. The default value is none.

post Specifies the mechanism that the monitor uses for posting. The default value is RespFormat=HTML.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at

sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

time-until-up

Specifies the amount of time, in seconds, after the first successful response before a node is marked up. A value of 0 (zero) causes a node to be marked up immediately after a valid response is received from the node. The default value is 0 (zero).

timeout

Specifies the number of seconds the target has in which to respond to the monitor request. The default value is 16 seconds.

If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Also, if the target responds with a RESET packet, the system immediately flags the target as down without waiting for the timeout interval to expire.

url Specifies the URL that the monitor uses. The default value is /scripts/f5Isapi.dll.

username

Specifies the user name if the monitored target requires authentication. The default value is none.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2017-04-05 Itm monitor wmi(1)

Itm nat-stats

NAME nat-stats - Shows Network address translation roll up statistics.

MODULE

cgnat afm

SYNTAX

Display NAT stats per roll-up-level using the following syntax.

DISPLAY

show nat-stats roll-up-level [...]

name [...] stat [...]

options:

all (where applicable)

recursive

DESCRIPTION

This command displays NAT stats per roll up level specified. Supported stat types for each roll up level may vary.

roll-up-levels

lsn-pool - Provide the stats for lsn-pool.

fw-nat-source-translation-object - Provide the stats for FWNAT source translation object.

translation-address - Provide the stats for a translation address. Applicable only at nat-stats level high.

global - Provide the cumulative stats at global level.

stat types

all - Display all applicable stats for this level.

total-addr - Provide the current total of translation addresses at each reporting level.

total-endpoints - Provide the current total of translation endpoints at each reporting level.

total-port-blks - Provide the current number of total port blocks at each reporting level.

active-addr - Provide the current number of translation addresses that are actively being used.

active-endpoints - Provide the current number of active translation endpoints.

in-use-port-blks - Provide the current number of port blocks in use at each reporting level, if that level includes an address used in a NAT PBA pool.

active-subscribers - Provide the current number of unique subscriber addresses that have an active translation. Applicable only at nat-stats level medium.

active-connections - Provide the current number of active connections at each reporting level.

cumulative-subscribers - Provide the total number of unique subscribers that have connected for each reporting level. Applicable only at nat-stats level medium.

cumulative-translations - Provide the cumulative number of translation mappings that have occurred.

peak-subscribers - Provide count of max number of unique, concurrent subscribers. Applicable only at nat-stats level medium.

peak-translations - Provide count of max number of concurrent translations.

port-reservations - Provide counts of number of port reservations by Application-Layer Gateways (ALGs).

hairpin-connections - Provide the cumulative number of hairpin connections at each reporting level.

persistence-entries - Provide the number of persistence entries. Applicable only at nat-stats level high.
inbound-entries - Provide counts of inbound entries, flows, and failures.
inbound-connections - Provide the number of active inbound connections.
inbound-failures - Provide the number of inbound failures.
failure-count - Provide the number of translation errors.

EXAMPLES

```
show ltm nat-stats roll-up-level fw-nat-source-translation-object name /Common/snat1 stat all Displays all available stats for roll up type source translation object /Common/snat1.
```

SEE ALSO

show, ltm nat

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2019-06-23 ltm nat-stats(1)

ltm nat

NAME

nat - Configures network address translation (NAT) for the Local Traffic Manager.

MODULE

ltm

SYNTAX

Configure the nat component within the ltm module using the syntax in the following sections.

CREATE/MODIFY

```
create nat [name]
```

```
modify nat [name]
```

options:

```
app-service [[string] | none]
```

```
arp
```

```
auto-lasthop [default | enabled | disabled ]
```

```
description [string]
```

```
[enabled | disabled]
```

```
originating-address [ip address]
```

```
translation-address [ip address]
```

```
traffic-group [[string] | default | non-default | none]
```

```
vlan [add | delete | replace-all-with] {
```

```
    [vlan names...]
```

```
}
```

```
vlan [default | none]
```

```
vlan-disabled
```

```
vlan-enabled
```

```
reset-stats nat
```

```
reset-stats nat [ [ [name] | [glob] | [regex] ] ... ]
```

```
edit nat [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list nat
```

```
list nat [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config nat
```

```
show running-config nat [ [ [name] | [glob] |
```

```
    [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

```
partition
```

```
show nat
```

```
show nat [name]
```

options:

```
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
```

```
field-fmt
```

DELETE

```
delete nat [name]
```

DESCRIPTION

A network address translation (NAT) defines a mapping between an originating IP address and an IP address that you specify.

A primary reason for defining a NAT is to allow one of the servers in the server array behind the traffic management system to start communication with a computer in front of, or external to, the system.

EXAMPLES

```
create nat new_nat translation-address 10.0.140.100 originating-address 11.0.0.100
```

The node behind the system with the IP address 10.0.140.100 has a presence in front of the BIG-IP(r) System as IP address 11.0.0.100.

```
delete nat new_nat
```

Permanently deletes the NAT from the system configuration.

Additional Restrictions

The nat component has the following additional restrictions:

- A virtual server cannot use the IP address specified in the NAT.
- A NAT should not use an IP address of a BIG-IP system.
- A NAT cannot use an originating or translated IP address defined for and used by a SNAT or another NAT.
- You must delete a NAT before you can redefine it.

OPTIONS

app-service

Specifies the name of the application service to which the NAT belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the NAT. Only the application service can modify or delete the NAT.

arp Enables or disables Address Resolution Protocol (ARP).

description

User defined description.

[enabled | disabled]

Enables or disables the NAT. The default value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

inherited-traffic-group

Indicates if the traffic-group is inherited from the parent folder. This property is read only.

originating-address

Specifies the IP address from which traffic is being initiated.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

traffic-group

Specifies the traffic group of the failover device group on which the NAT is active. The default traffic group is inherited from the containing folder.

translation-address

Specifies the IP address that is translated or mapped, and the IP address to which it is translated or mapped. This option is required when creating a NAT. This option may not be changed after the nat has been created.

unit Specifies the unit in a redundant system. Derived from traffic-group. This property is read only.

vlan

Specifies a list of existing VLANs on which access to the NAT is enabled or disabled. A NAT is accessible on all VLANs by default.

vlan-disabled

Indicates the NAT is disabled on the list of VLANs.

vlan-enabled

Indicates the NAT is enabled on the list of VLANs.

SEE ALSO

create, delete, edit, glob, list, ltm snat, ltm snat-translation, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

Itm node

NAME

node - Configures node addresses and services.

MODULE

itm

SYNTAX

Configure the node component within the ltm module using the syntax in the following sections.

CREATE/MODIFY

create node [name]

modify node [name]

options:

address [ip address]

app-service [[string] | none]

connection-limit [integer]

description [string]

[down | up]

dynamic-ratio [integer]

fqdn {

name [string]

address-family [ipv4 | ipv6]

autopopulate [enabled | disabled]

down-interval [integer]

interval [integer]

}

logging [enabled | disabled]

monitor [[name] | none]

rate-limit [integer]

ratio [integer]

session [user-enabled | user-disabled]

state [user-down | user-up]

metadata

[add | delete | modify] {

[metadata_name ...] {

value ["value content"]

persist [true | false]

}

}

reset-stats node

reset-stats node [[[ip address] | [glob] | [regex]] ...]

edit node [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv node [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list node

list node [[[name] | [glob] | [regex]] ...]

show running-config node

show running-config node [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show node

show node [name]

options:

all-properties

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete node [name]

DESCRIPTION

Displays information about nodes, and sets attributes of nodes and node IP addresses.

EXAMPLES

list node all-properties

Displays all of the properties of all of the nodes.

modify node all monitor none

Removes all monitor associations from nodes.

create node myNode address 10.10.10.15

Creates a node named myNode with an IP address of 10.10.10.15.

modify node myNode monitor none

Removes all monitor associations from the node, myNode.

show node

Displays statistics and status for all nodes in the system configuration.

show node all-properties

Displays statistics and status for all nodes in the system configuration. If the system includes Packet Velocity(r) ASIC (PVA) and PVA Assist capabilities, this command displays status and statistics for that feature.

mv /ltm node /Common/10.10.10.15 to-folder /Common/all_nodes

Moves the node 10.10.10.15 to a folder named all_nodes, where all_nodes has already been created under /Common.

Note: If you wish to change the name of the node, you must use the nodes same IP Address or a name that does not represent an IP Address that does not match the address configured on the node.

Please refer to the mv manual page for additional examples on how to use the mv command.

OPTIONS

address

Specifies the IP address of the node.

app-service

Specifies the name of the application service to which the node belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the node. Only the application service can modify or delete the node.

connection-limit

Specifies the maximum number of connections that a node or node address can handle. The default value is 0 (zero).

description

Specifies a user-defined description.

[down | up]

Marks the node up or down. The default value is down.

dynamic ratio

Sets the dynamic ratio number for the node. The ratio weights are based on continuous monitoring of the servers and are therefore continually changing. The default value is 1.

Dynamic Ratio load balancing can currently be implemented on RealNetworks RealServer platforms, on Windows platforms equipped with Windows Management Instrumentation (WMI), or on a server equipped with either the UC Davis SNMP agent or Windows 2000 Server SNMP agent.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

metadata

Associates user defined data, each of which has a name and value pair and persistence. The default value is persistent, which saves the data to the config file.

fqdn Specifies the attributes for defining a fully qualified domain name for the node.

name Specifies the fully-qualified domain name of the node.

address-family

Specifies whether the fqdn should consider IPv4, IPv6, or IP-agnostic address family.

autopopulate

Specifies whether a node defined by a fully-qualified domain name should automatically scale to the set of IP addresses returned by the DNS query. If disabled, only one ephemeral node is generated from the first IP address returned by DNS. The default is disabled.

interval

Specifies the interval, in seconds, to instantiate DNS queries on a fully-qualified domain name. The default is 3600. A value of 'ttl' uses the TTL value obtained from the DNS server.

down-interval
Specifies the interval for the domain name resolution operation when a DNS query fails. The default is 5.

logging
Specifies whether the monitor applied should log its actions. Logs are stored in `/var/log/monitors/` and are regularly rotated and compressed. The default value is disabled. This option isn't a part of configuration and will reset to disabled on load. This option doesn't sync.

monitor
Specifies the name of the monitor that you want to associate with the node. The default value is none.

partition
Displays the administrative partition in which the node object resides.

rate-limit
Specifies the maximum number of connections per second allowed for a node or node address. The default value is 'disabled'.

ratio
Specifies the fixed ratio value used for a node during Ratio load balancing. The default value is 1.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

session
Establishing a session with a node is establishing the ability of the client to persist to the node when making new connections. When a node is session disabled, clients that have already established sessions with the node may create new connections, but a client that has not already established a session may not create a new one (or make a connection which would create a new session). This feature is used to gently drain connections from a node, typically as part of a maintenance operation. The default value is user-enabled.

state
Specifies the current state of the node. Use `user-down` to indicate that the node may not handle any new connections. Use `user-up`, after using `user-down`, to indicate that the node may accept new connections.

to-folder
This is used with the `mv` command to specify a folder in which to move the node to.

Note: nodes can be moved to any folder under `/Common`, but dependencies upon it may restrict moving it out of `/Common`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm pool`, `modify`, `mv`, `regex`, `reset-stats`, `show`, `tmsb`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2015. All rights reserved.

BIG-IP 2018-10-17 ltm node(1)

ltm persistence cookie

NAME
cookie - Configures a cookie persistence profile.

MODULE
ltm persistence

SYNTAX
Configure the cookie component within the ltm persistence module using the syntax in the following sections.

MODIFY
create cookie [name]
modify cookie [name]
options:
all
always-send [enabled | disabled]
app-service [[string] | none]
cookie-name [[name] | none]
cookie-encryption [required | preferred | disabled]
cookie-encryption-passphrase [string | none]

```

defaults-from [name]
description [string]
expiration [ [d:h:m:s] | [h:m:s] | [m:s] | [seconds]
| "session cookie" ]
httponly [enabled | disabled]
secure [enabled | disabled]
hash-length [integer]
hash-offset [integer]
match-across-pools [enabled | disabled]
match-across-services [enabled | disabled]
match-across-virtuals [enabled | disabled]
method [hash | insert | passive | rewrite]
mirror [enabled | disabled]
override-connection-limit [enabled | disabled]
timeout [indefinite | [integer] ]
encrypt-cookie-poolname [enabled | disabled]

edit cookie [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

mv cookie [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
options:
  to-folder

DISPLAY
list cookie
list cookie [ [ [name] | [glob] | [regex] ] ... ]
show running-config cookie
show running-config cookie [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  non-default-properties
  one-line
  partition

DELETE
delete cookie [name]
options:
  all

```

DESCRIPTION

You can use the cookie component to configure cookie persistence for the BIG-IP(r) system. Cookie persistence uses an HTTP cookie stored on a client's computer to allow the client to connect to the same server previously visited at a web site.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server. Using a persistence profile avoids having to write an iRule to implement a type of persistence. You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

```
list cookie
```

Displays all cookie persistence profiles.

```
create cookie cookie_persistence defaults-from cookie
```

Creates a custom cookie persistence profile named `cookie_persistence` that inherits its settings from the default cookie persistence profile.

```
mv cookie /Common/my_cookie_profile to-folder /Common/my_folder
```

Moves a custom cookie persistence profile named `my_cookie_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`always-send`

Send the cookie persistence entry on every reply, even if the entry has previously been supplied to the client. The default value is disabled.

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`cookie-name`

Specifies a unique name for the cookie. This option is required.

`defaults-from`

Specifies the existing profile from which the system imports settings for the new profile. The default value is `cookie`, the system default cookie persistence profile.

`description`

User defined description.

cookie-encryption

Specifies the way in which cookie format will be used: disabled: generate old format, unencrypted, preferred: generate encrypted cookie but accept both encrypted and old format, and required: cookie format must be encrypted. Default is required.

cookie-encryption-passphrase

Specifies a passphrase to be used for cookie encryption.

expiration

Specifies the cookie expiration date in the format d:h:m:s, h:m:s, m:s or seconds. (hours 0-23, minutes 0-59, seconds 0-59). The time period must be less than 24856 days.

You can use "session-cookie" (0 seconds) to indicate that the cookie expires when the browser closes.

encrypt-cookie-poolname

Specifies whether the pool-name in the inserted BigIPServer default cookie should be encrypted. The default value is disabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

httponly

Specifies whether the httponly attribute should be enabled or disabled for the inserted cookies. The default value is enabled.

secure

Specifies whether the secure attribute should be enabled or disabled for the inserted cookies. The default value is enabled.

hash-length

Specifies the cookie hash length. The length is the number of bytes to use when calculating the hash value. The default value is 0 (zero) bytes.

hash-offset

Specifies the cookie hash offset. The offset is the number of bytes in the cookie to skip before calculating the hash value. The default value is 0(zero) bytes.

match-across-pools

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

match-across-services

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

match-across-virtuals

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

method

Specifies the type of cookie processing that the system uses. The default value is insert.

mirror

Specifies whether the system mirrors persistence records to the high-availability peer. This option is applicable only when the value of the method option is hash. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

override-connection-limit

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the duration of the persistence entries. The default value is 180 seconds.

to-folder

cookie persistence profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015-2016. All rights reserved.

ltm persistence dest-addr

NAME

dest-addr - Configures a destination address affinity persistence profile.

MODULE

ltm persistence

SYNTAX

Configure the dest-addr component within the ltm persistence module using the syntax in the following sections.

MODIFY

create dest-addr [name]

modify dest-addr [name]

options:

- all
- app-service [[string] | none]
- defaults-from [name]
- description [string]
- hash-algorithm [carp | default]
- mask [[ip address] | none]
- match-across-pools [enabled | disabled]
- match-across-services [enabled | disabled]
- match-across-virtuals [enabled | disabled]
- mirror [enabled | disabled]
- override-connection-limit [enabled | disabled]
- timeout [integer]

edit dest-addr [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties

mv dest-addr [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

- to-folder

DISPLAY

list dest-addr

list dest-addr [[[name] | [glob] | [regex]] ...]

show running-config dest-addr

show running-config dest-addr [[[name] | [glob] | [regex]] ...]

options:

- all
- all-properties
- non-default-properties
- one-line
- partition

DELETE

delete dest-addr [name]

options:

- all

DESCRIPTION

You can use the dest-addr component to configure a destination address affinity persistence profile for the BIG-IP(r) system. Also known as sticky persistence, destination address affinity persistence supports TCP and UDP protocols, and directs session requests to the same server based solely on the destination IP address of a packet.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server. Using a persistence profile means that you do not have to write an iRule to implement a type of persistence. You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

```
list dest-addr
```

Displays all destination address affinity persistence profiles.

```
create dest-addr da_persistence defaults-from dest-addr
```

Creates a custom destination address affinity persistence profile named da_persistence that inherits its settings from the default destination address affinity persistence profile.

```
mv dest-addr /Common/my_dest-addr_profile to-folder /Common/my_folder
```

Moves a custom destination address persistence profile named `my_dest-addr_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`defaults-from`

Specifies the existing profile from which the system imports settings for the new profile. The default value is `dest_addr`, the system default destination address affinity persistence profile.

`description`

User defined description.

`glob` Displays the items that match the `glob` expression. See help `glob` for a description of `glob` expression syntax.

`hash-algorithm`

Specifies the system uses hash persistence load balancing. The default value is `default` (no hash persistence).

The options are:

`carp` Specifies to use the Cache Array Routing Protocol (CARP) to select the pool member for LB. The input to CARP is the hash value of destination address.

`default`

no hash persistence.

`mask` Specifies an IP mask. This is the mask used by simple persistence for connections. The default value is `::`.

`match-across-pools`

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

`match-across-services`

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

`match-across-virtuals`

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

`mirror`

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`override-connection-limit`

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`timeout`

Specifies the duration of the persistence entries. The default value is 180 seconds.

`to-folder`

`dest-addr` persistence profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `mv`, `regex`, `show`, `tms`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

Itm persistence global-settings

NAME

global-settings - Configures persistence for the BIG-IP(r) system.

MODULE

itm persistence

SYNTAX

Configure the global-settings component within the Itm persistence module using the syntax in the following sections.

MODIFY

modify global-settings [option name]

options:

description [string]

dest-addr-limit-mode [timeout | maxcount]

dest-addr-max [integer]

proxy-group [string]

edit global-settings [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list global-settings

list global-settings [[[name] | [glob] | [regex]] ...]

show running-config global-settings

show running-config global-settings

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the global-settings component within the Itm persistence module to configure persistence for the system.

For information about configuring session persistence for a virtual server, see the man pages for the following components: Itm persistence hash, Itm persistence msrdp, Itm persistence sip, Itm persistence source-addr, Itm persistence ssl, and Itm persistence universal.

EXAMPLES

list global-settings

Displays the global persistence configuration.

modify global-settings dest-addr-limit-mode maxcount

Sets the value of the option dest-addr-limit-mode to maxcount, which indicates that a persistence session is limited by the maximum number of requests to the destination address.

OPTIONS

description

User defined description.

dest-addr-limit-mode

Specifies that a persistence session is limited by either the number of seconds before the persistence entry times out, or by a maximum number of requests to the destination address. The default value is timeout.

dest-addr-max

Specifies the maximum number of entries that the persistence table can contain at any one time, when the value of the option dest-addr-limit-mode is maxcount. The default value is 2048 entries.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

proxy-group

Specifies a group of servers that are configured to process all of the requests from a single source address during a persistence session. The default value is aol.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

list, Itm virtual, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013. All rights reserved.

BIG-IP 2013-04-12 Itm persistence global-settings(1)

Itm persistence hash

NAME

hash - Configures a hash persistence profile.

MODULE

Itm persistence

SYNTAX

Configure the hash component within the Itm persistence module using the syntax in the following sections.

MODIFY

create hash [name]

modify hash [name]

options:

all

app-service [[string] | none]

defaults-from [name]

description [string]

hash-algorithm [carp | default]

hash-buffer-limit [integer]

hash-end-pattern [none | [string]]

hash-length [integer]

hash-offset [integer]

hash-start-pattern [none | [string]]

match-across-pools [enabled | disabled]

match-across-services [enabled | disabled]

match-across-virtuals [enabled | disabled]

mirror [enabled | disabled]

override-connection-limit [enabled | disabled]

rule [iRule name]

timeout [integer]

edit hash [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv hash [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list hash

list hash [[[name] | [glob] | [regex]] ...]

show running-config hash

show running-config hash [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

non-default-properties

one-line

partition

DELETE

delete hash [name]

options:

all

DESCRIPTION

You can use the hash component to configure a hash persistence profile for the BIG-IP(r) system. Hash persistence can also be activated from an existing iRule.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server.

Using a persistence profile means that you do not have to write an iRule to implement a type of persistence.

You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

list hash

Displays all hash persistence profiles.

create hash hash_persistence defaults-from hash

Creates a custom hash persistence profile named hash_persistence that inherits its settings from the default hash persistence profile.

mv hash /Common/my_hash_profile to-folder /Common/my_folder

Moves a custom hash persistence profile named my_hash_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is hash, the system default cookie persistence profile.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

hash-algorithm

Specifies the algorithm the system uses for hash persistence load balancing. The default value is default.

The options are:

carp Specifies to use the Cache Array Routing Protocol (CARP) to select the pool member for LB.

default

Specifies to use indexing of pool members to select the pool member for LB.

hash-buffer-limit

Specifies the maximum buffer length the system collects to locate the hashing pattern for hash persistence load balancing. The default value is 0 (zero).

hash-end-pattern

Specifies the string that describes the ending location of the hash pattern that the system uses to perform hash persistence load balancing. The default value is none.

hash-length

Specifies the length of data within the packet in bytes that the system uses to calculate the hash value when performing hash persistence load balancing. The default value is 0 (zero) bytes.

hash-offset

Specifies the start offset within the packet from which the system begins the hash when performing hash persistence load balancing. The default value is 0 (zero).

hash-start-pattern

Specifies the string that describes the start location of the hash pattern that the system uses to perform hash persistence load balancing. The default value is none.

match-across-pools

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

match-across-services

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

match-across-virtuals

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

mirror

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

override-connection-limit

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at

sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rule Specifies a rule name, if you are using a rule for universal persistence.

timeout

Specifies the duration of the persistence entries. The default value is 180 seconds.

to-folder

hash persistence profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2014-01-14 ltm persistence hash(1)

Itm persistence host

NAME

host - Configures a host persistence profile.

MODULE

ltm persistence

SYNTAX

Configure the host component within the ltm persistence module using the syntax in the following sections.

MODIFY

create host [name]

modify host [name]

options:

all

app-service [[string] | none]

defaults-from [name]

description [string]

match-across-pools [enabled | disabled]

match-across-services [enabled | disabled]

match-across-virtuals [enabled | disabled]

mirror [enabled | disabled]

override-connection-limit [enabled | disabled]

timeout [integer]

edit host [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv host [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list host

list host [[name] | [glob] | [regex]] ...]

show running-config host

show running-config host [name]

options:

all

all-properties

non-default-properties

one-line

partition

DELETE

delete host [name]

options:

all

DESCRIPTION

You can use the host component to configure a host persistence profile for the BIG-IP(r) system. Host persistence uses the HTTP Host header passed in a HTTP request to determine which pool member to pick. Host

persistence can also be activated from an existing iRule.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server. Using a persistence profile means that you do not have to write an iRule to implement a type of persistence. You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

list host

Displays all host persistence profiles.

create host host_persistence defaults-from host

Creates a custom host persistence profile named host_persistence that inherits its settings from the default host persistence profile.

mv host /Common/my_host_profile to-folder /Common/my_folder

Moves a custom host persistence profile named my_host_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is host, the system default host persistence profile.

description

User defined description.

match-across-pools

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

match-across-services

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

match-across-virtuals

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

mirror

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

override-connection-limit

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the duration of the persistence entries. The default value is 180 seconds.

to-folder

host persistence profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

Itm persistence msrdp

NAME

msrdp - Configures a Microsoft(r) Remote Display Protocol (MSRDP) persistence profile.

MODULE

itm persistence

SYNTAX

Configure the msrdp component within the Itm persistence module using the syntax in the following sections.

MODIFY

create msrdp [name]

modify msrdp [name]

options:

- all
- app-service [[string] | none]
- defaults-from [name]
- description [string]
- has-session-dir [no | yes]
- match-across-pools [enabled | disabled]
- match-across-services [enabled | disabled]
- match-across-virtuals [enabled | disabled]
- mirror [enabled | disabled]
- override-connection-limit [enabled | disabled]
- timeout [integer]

edit msrdp [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties

mv msrdp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

- to-folder

DISPLAY

list msrdp

list msrdp [[[name] | [glob] | [regex]] ...]

show running-config msrdp

show running-config msrdp [[[name] | [glob] | [regex]] ...]

options:

- all
- all-properties
- non-default-properties
- one-line
- partition

DELETE

delete msrdp [name]

options:

- all

DESCRIPTION

You can use the msrdp component to configure an MSRDP persistence profile for the BIG-IP(r) system. MSRDP persistence provides an efficient way of load balancing traffic and maintaining persistent sessions between Windows clients and servers that are running the Microsoft Terminal Services service. The recommended scenario for enabling the MSRDP persistence feature is to create a load balancing pool that consists of members running Windows .NET Server 2003, Enterprise Edition, or later, where all members belong to a Windows cluster and participate in a Windows session directory.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server. Using a persistence profile means that you do not have write an iRule to implement a type of persistence. You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

```
list msrdp
```

Displays all MSRDP persistence profiles.

```
create msrdp msrdp_persistence defaults-from msrdp
```

Creates a custom MSRDP persistence profile named msrdp_persistence that inherits its settings from the default MSRDP persistence profile

```
mv msrdp /Common/my_msrdp_profile to-folder /Common/my_folder
```

Moves a custom MSRDP persistence profile named my_msrdp_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is msrdp, the system default cookie persistence profile.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

has-session-dir

Specifies whether the Microsoft terminal services are configured with a session directory, so the system does not load balance the initial connection. The default value is yes.

match-across-pools

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

match-across-services

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

match-across-virtuals

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

mirror

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

override-connection-limit

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the duration of the persistence entries. The default value is 300 seconds.

to-folder

msrdp persistence profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2014-01-14 ltm persistence msrdp(1)

ltm persistence persist-records

NAME

persist-records - Displays or deletes persistence records.

MODULE

ltm persistence

SYNTAX

Configure the persist-records component within the ltm persistence module using the syntax in the following sections.

DISPLAY

show persist-records

options:

client-addr [ip address]
key [string]
mode [cookie | destination-address | hash | msrdp | sip |
source-address | ssl-session-id | universal]
node-addr [ip address]
node-port [integer]
pool [string]
save-to-file [filename]
virtual [string]

DELETE

delete persist-records

options:

client-addr [ip address]
key [string]
mode [cookie | destination-address | hash | msrdp | sip |
source-address | ssl-session-id | universal]
node-addr [ip address]
node-port [integer]
pool [string]
virtual [string]

DESCRIPTION

You can use the persist-records component to either display or delete records of persistent connections.

EXAMPLES

show persist-records

Displays all persistent connections on the BIG-IP(r) system.

delete persist-records client-addr 172.19.255.1

Deletes all persistent connections that originate from the client IP address, 172.19.255.1.

OPTIONS

client-addr

Specifies the IP address of the client from which the persistent connections you want to view or delete persist.

key Specifies a string that the system is using to persist the connections you want to view or delete.

mode Specifies the type of persistence of the connections you want to view or delete. The options are:

cookie

Cookie persistence uses an HTTP cookie stored on a client's computer to allow the client to connect to the same server previously visited at a web site.

destination-address

Also known as sticky persistence, destination address affinity persistence supports TCP and UDP protocols, and directs session requests to the same server based solely on the destination IP address of a packet.

hash Hash persistence is based on an existing iRule.

msrdp

MSRDP persistence provides an efficient way of load balancing traffic and maintaining persistent sessions between Windows(r) clients and servers that are running the Microsoft(r) Terminal Services service. The recommended scenario for enabling the MSRDP persistence feature is to create a load balancing pool that consists of members running Windows .NET Server 2003, Enterprise Edition, or later, where all members belong to a Windows cluster and participate in a Windows session directory.

sip Session Initiation Protocol (SIP) persistence is a type of persistence available for server pools.

You can configure SIP persistence for proxy servers that receive SIP messages sent through UDP. The BIG-IP system currently supports persistence for SIP messages sent through UDP, TCP, or SCTP.

source-address

Also known as simple persistence, source address affinity persistence supports TCP and UDP protocols, and directs session requests to the same server based solely on the source IP address of a packet. When you specify a source address as the mode of persistence, you must specify an IP address using the client-addr option.

ssl-session-id

SSL persistence is a type of persistence that tracks non-terminated SSL sessions, using the SSL session ID. Even when the client's IP address changes, the system still recognizes the connection as being persistent based on the session ID. Note that the term, non-terminated SSL sessions, refers to sessions in which the system does not perform the tasks of SSL certificate authentication and encryption/re-encryption.

universal

Universal persistence allows you to write an expression that defines what to persist on in a packet. The expression, written using the same expression syntax that you use in iRules(r), defines some sequence of bytes to use as a session identifier.

node-addr

Specifies the IP address of the node with which the client sessions that you want to view or delete remain persistent.

`node-port`
Specifies the port number of the node with which the client sessions that you want to view or delete remain persistent.

`pool` Specifies the pool member with which the client sessions that you want to view or delete remain persistent.

`save-to-file`
Specifies the file which persist-records information can be save to. With this option, it can write a file larger than 2GB.

`virtual`
Specifies the virtual server with which the client sessions that you want to view or delete remain persistent.

SEE ALSO

`delete`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013. All rights reserved.

BIG-IP 2013-04-10 ltm persistence persist-records(1)

ltm persistence sip

NAME

`sip` - Configures a Session Initiation Protocol (SIP) persistence profile.

MODULE

ltm persistence

SYNTAX

Configure the `sip` component within the ltm persistence module using the syntax in the following sections.

MODIFY

```
create sip [name]
modify sip [name]
options:
  all
  app-service [[string] | none]
  defaults-from [name]
  description [string]
  match-across-pools [ enabled | disabled]
  match-across-services [enabled | disabled]
  match-across-virtuals [enabled | disabled]
  mirror [enabled | disabled]
  override-connection-limit [enabled | disabled]
  sip-info [Call-ID | From | none | SIP-ETag | Subject | To]
  timeout [integer]
```

```
edit sip [ [ name ] | [ glob ] | [ regex ] ] ... ]
options:
  all-properties
  non-default-properties
```

```
mv sip [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
options:
  to-folder
```

DISPLAY

```
list sip
list sip [ [ name ] | [ glob ] | [ regex ] ] ... ]
show running-config sip
show running-config sip [ [ name ] | [ glob ] | [ regex ] ] ... ]
options:
  all
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

delete sip [name]
options:
all

DESCRIPTION

You can use the sip component to configure a SIP persistence profile for the BIG-IP(r) system. SIP persistence is a type of persistence available for server pools. You can configure SIP persistence for proxy servers that receive SIP messages sent through UDP. The BIG-IP system currently supports persistence for SIP messages sent through UDP, TCP, or SCTP.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server. Using a persistence profile means that you do not have to write an iRule to implement a type of persistence. You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

list sip

Displays all SIP persistence profiles.

create sip sip_persistence defaults-from sip_info

Creates a custom SIP persistence profile named sip_persistence that inherits its settings from the default SIP persistence profile.

mv sip /Common/my_sip_profile to-folder /Common/my_folder

Moves a custom SIP persistence profile named my_sip_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is sip_info, the system default cookie persistence profile.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

match-across-pools

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

match-across-services

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

match-across-virtuals

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

mirror

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

override-connection-limit

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

sip-info

Specifies the SIP header field on which you want SIP sessions to persist. The default value is none.

timeout

Specifies the duration of the persistence entries. The default value is 180 seconds.

to-folder

sip persistence profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2014-01-14 ltm persistence sip(1)

ltm persistence source-addr

NAME

source-addr - Configures a source address affinity persistence profile.

MODULE

ltm persistence

SYNTAX

Configure the source-addr component within the ltm persistence module using the syntax in the following sections.

MODIFY

create source-addr [name]

modify source-addr [name]

options:

all
app-service [[string] | none]
defaults-from [name]
description [string]
map-proxies [enabled | disabled]
map-proxy-address [ip address]
map-proxy-class [class name]
hash-algorithm [carp | default]
mask [[ip address] | none]
match-across-pools [enabled | disabled]
match-across-services [enabled | disabled]
match-across-virtuals [enabled | disabled]
mirror [enabled | disabled]
override-connection-limit [enabled | disabled]
timeout [integer]

edit source-addr [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

mv source-addr [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list source-addr

list source-addr [[[name] | [glob] | [regex]] ...]

show running-config source-addr

show running-config source-addr [[[name] | [glob] | [regex]] ...]

options:

all
all-properties
non-default-properties
one-line
partition

DELETE

delete source-addr [name]

options:

all

DESCRIPTION

You can use the source-addr component to configure a source address affinity persistence profile for the BIG-IP(r) system. Also known as simple persistence, source address affinity persistence supports TCP and UDP protocols, and directs session requests to the same server based solely on the source IP address of a packet.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server. Using a persistence profile means that you do not have to write an iRule to implement a type of persistence. You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

list source-addr

Displays all source address affinity persistence profiles.

create source-addr simple_persistence defaults-from source_addr

Creates a custom source address affinity persistence profile named simple_persistence that inherits its settings from the default source address affinity persistence profile.

mv source-addr /Common/my_source-addr_profile to-folder /Common/my_folder

Moves a custom source address persistence profile named my_source-addr_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is source_addr, the system default cookie persistence profile.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

hash-algorithm

Specifies the system uses hash persistence load balancing. The default value is default (no hash persistence).

The options are:

carp Specifies to use the Cache Array Routing Protocol (CARP) to select the pool member for LB. The input to CARP is the hash value of source address.

default

no hash persistence.

map-proxies

Enables or disables the map proxies attribute. The default value is disabled.

This attribute controls whether a source address will first be checked against an IP data-group/class to determine whether it is a well-known proxy address. If it matches the IP class, then the source address will be mapped to a single IP address for the purposes of persistence. The default well known proxy class is based on a pre-defined data-group "aol" which represents the AOL(r) company's previously-published list of proxies. Using this feature enables you to use client/source IP address persistence with a simple persist mask, but forces all clients matching the IP class to persist to the same server. The IP data-group/class may also be changed using either the map-proxy-class profile attribute or globally by changing the DB variable Persist.WellKnownProxyClass. Also, the IP address used for mapping a single source IP address for persistence may also be specifically set using the map-proxy-address profile attribute.

map-proxy-address

Specifies the single IP address to use when the source address matches the proxy data-group/class. The default value is any which results in the default behavior of using one of the IP addresses from the proxy data-group/class. Note: This mapped IP address does not have to be contained in the IP data-group/class. It may actually be any IP address since it is only used for keying the persistence record.

map-proxy-class

Specifies the data-group/class to use for determining whether a source address is from a proxy. The default value is none which will result in map_proxies using the class defined by the DB variable Persist.WellKnownProxyClass.

mask Specifies an IP mask. This is the mask used by simple persistence for connections. The default value is ::.

match-across-pools

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

match-across-services

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

match-across-virtuals

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

mirror

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

override-connection-limit
Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

partition
Displays the administrative partition within which the component resides.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout
Specifies the duration of the persistence entries. The default value is 180 seconds.

to-folder
source-addr persistence profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO
create, delete, edit, glob, list, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm persistence source-addr(1)

ltm persistence ssl

NAME
ssl - Configures a Secure Socket Layer (SSL) persistence profile.

MODULE
ltm persistence

SYNTAX
Configure the ssl component within the ltm persistence module using the syntax in the following sections.

MODIFY
create ssl [name]
modify ssl [name]
options:
all
app-service [[string] | none]
defaults-from [name]
description [string]
match-across-pools [enabled | disabled]
match-across-services [enabled | disabled]
match-across-virtuals [enabled | disabled]
mirror [enabled | disabled]
override-connection-limit [enabled | disabled]
timeout [integer]

edit ssl [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

mv ssl [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]
options:
to-folder

DISPLAY
list ssl
list ssl [[[name] | [glob] | [regex]] ...]
show running-config ssl
show running-config ssl [[[name] | [glob] | [regex]] ...]
options:
all
all-properties
non-default-properties
one-line

partition

DELETE

delete ssl [name]

options:

all

DESCRIPTION

You can use the ssl component to configure a destination address affinity persistence profile for the BIG-IP(r) system. SSL persistence is a type of persistence that tracks non-terminated SSL sessions, using the SSL session ID. Even when the client's IP address changes, the system still recognizes the connection as being persistent based on the session ID. Note that the term, non-terminated SSL sessions, refers to sessions in which the system does not perform the tasks of SSL certificate authentication and encryption/re-encryption.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server. Using a persistence profile means that you do not have to write an iRule to implement a type of persistence. You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

list ssl

Displays all SSL persistence profiles.

create ssl ssl_persistence defaults-from ssl

Creates a custom SSL persistence profile named ssl_persistence that inherits its settings from the default SSL persistence profile.

mv ssl /Common/my_ssl_profile to-folder /Common/my_folder

Moves a custom SSL persistence profile named my_ssl_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is ssl, the system default cookie persistence profile.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

match-across-pools

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

match-across-services

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

match-across-virtuals

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

mirror

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

override-connection-limit

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

timeout

Specifies the duration of the persistence entries. The default value is 300 seconds.

to-folder

ssl persistence profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2014-01-14 ltm persistence ssl(1)

ltm persistence universal

NAME

universal - Configures a universal persistence profile.

MODULE

ltm persistence

SYNTAX

Configure the universal component within the ltm persistence module using the syntax in the following sections.

MODIFY

create universal [name]

modify universal [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

match-across-pools [enabled | disabled]

match-across-services [enabled | disabled]

match-across-virtuals [enabled | disabled]

method [hash | insert | passive | rewrite]

mirror [enabled | disabled]

override-connection-limit [enabled | disabled]

rule [[iRule name] | none]

timeout [integer]

edit universal [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv universal [[source-name] [destination-name]] [[name] to-folder [folder-name]] [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list universal

list universal [[name] | [glob] | [regex]] ...]

show running-config universal

show running-config universal [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete universal [name]

DESCRIPTION

You can use the universal component to configure a universal persistence profile for the BIG-IP(r) system.

With universal persistence you can write an expression that defines what to persist on in a packet. The expression, written using the same expression syntax that you use in iRules(r), defines some sequence of bytes to use as a session identifier.

A persistence profile is a profile that enables persistence when you assign the profile to a virtual server.

Using a persistence profile means that you do not have to write an iRule to implement a type of persistence.

You can either use the default profile, or create a custom profile based on the default.

EXAMPLES

list universal

Displays all universal persistence profiles.

create universal uni_persistence defaults-from universal

Creates a custom universal persistence profile named `uni_persistence` that inherits its settings from the default universal persistence profile.

```
mv universal /Common/my_universal_profile to-folder /Common/my_folder
```

Moves a custom universal persistence profile named `my_universal_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`defaults-from`

Specifies the existing profile from which the system imports settings for the new profile. The default value is `universal`, the system default cookie persistence profile.

`description`

User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`match-across-pools`

Specifies, when enabled, that the system can use any pool that contains this persistence record. The default value is disabled.

`match-across-services`

Specifies, when enabled, that all persistent connections from a client IP address, which go to the same virtual IP address, also go to the same node. The default value is disabled.

`match-across-virtuals`

Specifies, when enabled, that all persistent connections from the same client IP address go to the same node. The default value is disabled.

`mirror`

Specifies whether the system mirrors persistence records to the high-availability peer. The default value is disabled.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`override-connection-limit`

Specifies, when enabled, that the pool member connection limits are not enforced for persisted clients. Per-virtual connection limits remain hard limits and are not disabled. The default value is disabled.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`rule` Specifies an iRule name when you are using a rule for universal persistence.

`timeout`

Specifies the duration of the persistence entries. The default value is 180 seconds.

`to-folder`

universal persistence profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `itm virtual`, `modify`, `mv`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2014-06-04 itm persistence universal(1)

NAME
policy-strategy - Configures policy-strategy for Centralized Policy Manager.

MODULE
ltm

DESCRIPTION

The policy-strategy component stores the different matching strategies employed by LTM Policy engine. Strategy comes into play when a policy has multiple rules, and the behavior of the policy can be customized as the situation requires.

There are 3 pre-defined matching strategies: "first-match", "all-match", and "best-match". A "first-match" strategy terminates the matching engine on the first condition that matches and executes that rule's actions. An "all-match" strategy will execute the actions for all conditions that match.

The "best-match" strategy is intended for situations when multiple conditions match simultaneously, and allows for the more specific match to win. For example, one rule may match the http-uri hostname while another may match the http-uri extension. The system has a built-in table defining combinations of event, operand, and selector, and an associated precedence value for each combination. When multiple rules match in a "best-match" situation, then the condition with the lowest ordinal value of event-operand-selector precedence is declared to be the most specific, and its actions are executed.

Generally policy-strategy should not require additions or changes. However, it could make sense to create user-defined policy-strategy when a "best-match" strategy is desired, but the built-in precedence table does not reflect the organization's idea of which operand-selector combinations are most specific.

For additional details, refer to Local Traffic Policy documentation on the AskF5 knowledge base at <http://support.f5.com>.

CREATE/MODIFY

create policy-strategy [name]
modify policy-strategy [name]
options:

[strategy | [all-match | best-match | first-match]]

operands [add | delete | modify | replace-all-with] {
ORDINAL {
[OPERAND] [EVENT] [SELECTOR]
}
}

[app-service [VALUE | none]]
[partition VALUE]

where

strategy
Specifies the match method: all-match, best-match, or first-match.

operands
Define a combination of event, operand, selector, and associate it with an ordinal precedence value.

ORDINAL
Integer precedence value, lower value indicates a higher precedence.

OPERAND
Entity to compare, see some examples in Precedence Table below, or ltm_policy documentation for list with descriptions.

EVENT
Framework event like request or response, default is "request" if not specified.

SELECTOR
More specific part of operand, default is "all" if not specified. See some examples in the Precedence Table below, or ltm_policy documentation for list and descriptions.

app-service
Specifies the name of the application service to which the policy strategy belongs. The default value is "none" if not specified. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the policy strategy. Only the application service can modify or delete the policy strategy.

DISPLAY

list policy-strategy
list policy-strategy [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE

delete policy-strategy [name]

EXAMPLES

create policy-strategy my_strategy strategy first-match

Creates the policy strategy my_strategy which matches the first rule selected.

Precedence Table

Ordinal Event Operand Selector

1	request	tcp	port
2	request	tcp	vlan-id
3	request	tcp	vlan
4	request	tcp	route-domain
5	request	tcp	rtt
6	request	tcp	mss
7	request	client-ssl	cipher
8	request	client-ssl	cipher-bits
9	request	http-host	host
10	request	http-host	port
11	request	http-host	all
12	request	http-version	all
13	request	http-version	major
14	request	http-version	minor
15	request	http-method	all
16	request	http-uri	scheme
17	request	http-uri	host
18	request	http-uri	port
19	request	http-uri	path-segment
20	request	http-uri	extension
21	request	http-uri	path
22	request	http-uri	query-parameter
23	request	http-uri	unnamed-query-parameter
24	request	http-uri	query-string
25	request	http-uri	all
26	request	http-cookie	all
27	request	http-basic-auth	username
28	request	http-basic-auth	password
29	request	http-referer	all
30	request	http-referer	scheme
31	request	http-referer	host
32	request	http-referer	port
33	request	http-referer	path-segment
34	request	http-referer	path
35	request	http-referer	extension
36	request	http-referer	query-parameter
37	request	http-referer	unnamed-query-parameter
38	request	http-referer	query-string
39	request	http-header	all
40	response	http-version	all
41	response	http-version	major
42	response	http-version	minor
43	response	http-status	all
44	response	http-status	code
45	response	http-status	text
46	response	http-header	all
47	request	geoip	org
48	request	geoip	isp
49	request	geoip	region-code
50	request	geoip	region-name
51	request	geoip	country-code
52	request	geoip	country-name
53	request	geoip	continent
54	request	cpu-usage	last-15secs
55	request	cpu-usage	last-1min
56	request	cpu-usage	last-5mins
57	request	http-user-agent	device-make
58	request	http-user-agent	device-model
59	request	http-user-agent	browser-type
60	request	http-user-agent	browser-version
61	request	http-user-agent	user-agent-token

SEE ALSO

ltm policy, create, delete, edit, glob, list, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

ltm policy

NAME

policy - Configures a policy for Centralized Policy Manager.

MODULE

ltm

SYNTAX

Create or modify LTM Policies within the ltm module, using the syntax shown in the following sections.

Policies exist in 2 forms, draft and published. Only draft policies can be modified, and only published policies can be applied to a virtual server. A draft policy can be turned into a published policy using the publish command. A draft copy can be obtained from a published policy using modify [name] create-draft. Draft policies are placed in a Drafts folder.

For additional details, refer to Local Traffic Policy documentation on the AskF5 knowledge base at <http://support.f5.com>.

CREATE/MODIFY

```
create policy Drafts/[name]
```

```
modify policy Drafts/[name]
```

options:

```
strategy [STRING | none]
```

```
copy-from [name | Drafts/name]
```

```
create-draft
```

```
rules [add | delete | modify | replace-all-with] {
```

```
STRING {
```

```
ordinal NUMBER |
```

```
app-service STRING |
```

```
conditions [add | delete | modify | replace-all-with] {
```

```
NUMBER { CONDITION_SPEC [[CONDITION_SPEC] ...] }
```

```
} |
```

```
actions [add | delete | modify | replace-all-with] {
```

```
NUMBER { ACTION_SPEC [[ACTION_SPEC] ...] }
```

```
}
```

```
}
```

```
}
```

```
[controls [add | delete | modify | replace-all-with] {
```

```
CONTROLS_ASPECT [[CONTROLS_ASPECT] ...]
```

```
}]
```

```
[requires [add | delete | modify | replace-all-with] {
```

```
REQUIRES_ASPECT [[REQUIRES_ASPECT] ...]
```

```
}]
```

PUBLISH

```
publish policy name
```

DISPLAY

```
list policy
```

```
list policy [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

```
partition
```

```
show policy
```

```
show policy [name]
```

options:

```
all-properties
```

```
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
```

```
detail
```

```
field-fmt
```

DELETE

```
delete policy [name]
```

Note: Before a policy can be deleted, it must be removed from all virtual servers holding a reference to it.

DESCRIPTION

An LTM Policy is a set of rules which can be attached to a virtual server to efficiently process traffic.

Similar in concept to iRules, Policies can inspect requests and responses, and perform programmed actions.

The controls and requires aspects for a policy are automatically set by the system based on an inspection of the conditions and actions specified in LTM Policy rules. User should not specify either of these.

EXAMPLES

```
create policy Drafts/my_policy
```

```
strategy my_strategy
```

Creates a Local Traffic Manager policy in the Drafts folder named my_policy. The strategy determining policy actions is my_strategy. Draft policies may be modified, but cannot be applied to a virtual server until they

are published.

create ltm policy Drafts/new_policy copy-from published_policy
create ltm policy Drafts/new_policy copy-from Drafts/old_policy

Creates a Local Traffic Manager policy based on an existing published policy, and from an existing draft policy, respectively.

modify ltm policy new_policy create-draft

Creates a draft policy of an existing published policy.

publish ltm policy Drafts/my_policy

Takes a policy that was created or modified in the Drafts folder, and publishes it. Published policies can then be applied to a virtual server.

delete policy my_policy

Deletes the policy named my_policy.

show policy

Displays statistics and status for all Local Traffic Manager policies in the system configuration.

show policy all-properties

Displays statistics and status for all Local Traffic Manager policies in the system configuration.

Note that if the system includes Packet Velocity(r) ASIC (PVA) and PVA Assist capabilities, this command displays status and statistics for that feature.

list policy my_policy

Displays properties of the policy named my_policy.

app-service - Specifies the name of the application service to which the policy belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the policy. Only the application service can modify or delete the policy.

CONTROLS_ASPECT is one or more of the following:

none - Associated with unrestricted actions that are always available, like logging.

forwarding - Enables many options for Forwarding connections to the back end.

caching - Caching enable or disable on a per-request basis.

compression - Actions which enable / disable compression and decompression. Virtual server will need to have an appropriate compression profile applied.

acceleration - Enable or disable Web Acceleration

asm - Enable or disable Application Security Management

avr - Enable or disable Application Visibility Reporting

l7dos - Enable or disable Layer 7 Denial-of-Service protection

bot-defense - Enable or disable Unified Bot Defense protection

classification - Used by the Traffic classification engine

request-adaptation - Enable or disable Request Adaptation

response-adaptation - Enable or disable Response Adaptation

client-ssl - Enable or disable SSL connection on the client side

server-ssl - Enable or disable SSL connection on the server side

websocket - Actions related to WebSockets

REQUIRES_ASPECT is one or more of the following:

none - Associated with conditions that are always available, like cpu-usage, with no specific profiles required to be attached to a virtual server

http - Makes available HTTP-protocol conditions. A profile that communicates using the HTTP protocol needs to be attached to a virtual server with this policy.

http-explicit - Makes available HTTP Explicit Proxy specific conditions. An HTTP Explicit Proxy profile needs to be attached to a virtual server with this policy.

http-connect - Makes available HTTP Connect specific conditions. An HTTP Connect profile needs to be attached to a virtual server with this policy.

tcp - A TCP profile needs to be attached to a virtual with this policy. Makes available TCP-specific conditions.

client-ssl - A Client-SSL profile needs to be attached to a virtual with this policy.

server-ssl - A Server-SSL profile needs to be attached to a virtual with this policy.

classification - A classification profile needs to be attached to a virtual with this policy.

Data Types

BOOLEAN - [true* | false]

NUMBER - signed 32-bit integer

STRING - Bare_string -or- "quoted string"

TCLSTRING - A STRING optionally containing Tcl command substitutions to be evaluated at runtime. If string begins with the 4-character prefix "tcl:", then the prefix is removed and the rest of the string is passed to the Tcl interpreter. If no prefix, then whole string is treated as a plain string with no Tcl interpreter overhead. Examples:

```
log request message "tcl:This is Tcl-enabled and the URI is [HTTP::uri]"
```

```
log request message "This is just a plain old string"
```

IP_ADDRESS - IPv4 or IPv6 address

Comparison operators

Core to defining conditions is the need to compare quantities at run time against pre-defined values. LTM Policy allows you to specify single or multiple values in a comparison.

NUMBER_COMPARISON

```
[not] [equals* | less | greater | less-or-equal | greater-or-equal]
values { VAL1 [[VAL2] ...]}
```

Sample numeric comparisons:

```
not greater values { 1024 } # no more than 1024
```

```
values { 80 443 8080 } # compare against 80, 443, or 8080, equals implied
```

STRING_COMPARISON

```
[not] [equals* | starts-with | ends-with | contains]
values { VAL1 [[VAL2] ...]} [case-insensitive* | case-sensitive ]
```

Sample string comparisons:

```
equals values { Abel bAkEr chArllE } # case-insensitive match of 3 candidates
```

```
values { Abel bAkEr chArllE } # same as above, equals implied
```

```
ends-with values { html txt } # match if string ends with either candidate
```

```
contains values { "jj83Q@@#AFRT@==" } case-sensitive # match value, case must match
```

*default if not specified

CONDITION_SPEC

A CONDITION_SPEC, or condition specification, is where you can tell the system the specific attributes you would like to inspect and use as a trigger for action.

Conditions are associated with an event, so conditions can be evaluated at different times during a request-response cycle.

Below is a list of all supported conditions, the events during which they can be evaluated, additional qualifiers, and parameters.

http-uri

Inspect the URI on a request and match on various parts or the entire URI (since 11.4.0)

Specifying http-uri in a condition automatically adds "requires {http}" to the policy.

```
http-uri [proxy-request* | request | proxy-connect]
[all STRING_COMPARISON
 [normalized BOOL]
]
[scheme STRING_COMPARISON
 [normalized BOOL]
]
[host STRING_COMPARISON
 [normalized BOOL]
]
[port NUMBER_COMPARISON
 [normalized BOOL]
]
[path STRING_COMPARISON
 [normalized BOOL]
]
```

```

[extension STRING_COMPARISON
 [normalized BOOL]
]
[query-string STRING_COMPARISON
 [normalized BOOL]
]
[query-param STRING_COMPARISON
 name STRING
 [normalized BOOL]
]
[unnamed-query-param STRING_COMPARISON
 index NUMBER
 [normalized BOOL]
]
[path-segment STRING_COMPARISON
 index NUMBER
 [normalized BOOL]
]
[urlcat
 [normalized BOOL]
]

```

where

all - match on the full URI
normalized - Convert URI to standard form for consistent comparison.
scheme - match on the scheme, e.g. http, https, ftp,
file
normalized - Convert URI to standard form for consistent comparison.
host - match on the hostname in the URI
normalized - Convert URI to standard form for consistent comparison.
port - match on the port number in the URI
normalized - Convert URI to standard form for consistent comparison.
path - match on the URI path
normalized - Convert URI to standard form for consistent comparison.
extension - match on the file extension in the URI, e.g. jpg, html, cgi
normalized - Convert URI to standard form for consistent comparison.
query-string - match against text in the query string
normalized - Convert URI to standard form for consistent comparison.
query-param - match value of the named query parameter from the query string
name - Specify the name of the particular query parameter whose value is to be used
normalized - Convert URI to standard form for consistent comparison.
unnamed-query-param - match the value of a query parameter by a numeric index instead of by name
index - The numeric order of the query parameter whose value is to be used, starting at 1. Negative values indicate counting right to left.
normalized - Convert URI to standard form for consistent comparison.
path-segment - Match a part of the URI path by a numeric index
index - The numeric order of a segment in the path, starting at 1. Negative values indicate counting right to left.
normalized - Convert URI to standard form for consistent comparison.
urlcat - Run URI through a categorization engine. List of categories - 'tmsh list sys url-db url-category'
normalized - Convert URI to standard form for consistent comparison.

tcp

Inspect and match on various TCP properties of a connection (since 11.5.0)

Specifying tcp in a condition automatically adds "requires {tcp}" to the policy.

```

tcp [client-accepted* | ssl-client-hello |
ssl-client-serverhello-send | ssl-server-hello |
ssl-server-handshake | server-connected |
request | response | proxy-request | proxy-connect |
proxy-response | ws_request | ws_response |
classification-detected]
 [address IP_COMPARISON
 [internal BOOL]
 [local BOOL]

```

```

]
[port NUMBER_COMPARISON
 [internal BOOL]
 [local BOOL]
]
[mss NUMBER_COMPARISON
 [internal BOOL]
]
[rtt NUMBER_COMPARISON
 [internal BOOL]
]
[vlan STRING_COMPARISON
 [internal BOOL]
]
[vlan-id NUMBER_COMPARISON
 [internal BOOL]
]
[route-domain NUMBER_COMPARISON
 [internal BOOL]
]

```

where

address - Match on IP address. By default the IP address is the one associated with the external interface, remote end of the connection.

internal - Internal specifies the IP address of the endpoint on the "internal" side of the connection.

local - Local specifies the IP address of the local side of the connection, i.e. not the remote side.

port - Match on port number. By default the port is the one associated with the external interface, remote end of the connection.

internal - Internal specifies the port of the endpoint on the "internal" side of the connection.

local - Local specifies the port of the local side of the connection, i.e. not the remote side.

mss - Compare the TCP maximum segment size on the external network interface.

internal - Refers to the maximum segment size on the internal interface.

rtt - Inspect the round trip time on the external network interface.

internal - Refers to the round trip time on the internal interface.

vlan - Compare traffic with specified vlan on the external network interface.

internal - Refers to the vlan on the internal interface.

vlan-id - Compare traffic with specified vlan-id number on the external network interface.

internal - Refers to the vlan-id on the internal interface.

route-domain - Compare traffic with specified route domain number on the external network interface.

internal - Specifies the route domain on the internal interface.

Examples

```

tcp address matches values { 141.202.53.16 }
tcp address internal matches values { 192.168.63.1 192.168.63.121 }
tcp port matches values { 8080 3128 }

```

client-ssl

Inspect properties of the SSL connection on the client side of the device. (since 11.4.0)

Specifying client-ssl in a condition automatically adds "requires {client-ssl}" to the policy.

```

client-ssl [proxy-request* | request | proxy-connect |
proxy-response | response]
 [protocol STRING_COMPARISON]
 [cipher STRING_COMPARISON]
 [cipher-bits NUMBER_COMPARISON]

```

where

protocol - Compare SSL protocol name

cipher - Cipher name

cipher-bits - cipher strength in number of bits

http-method

Inspect the request's HTTP method, e.g. GET, POST, HEAD (since 11.4.0)

Specifying http-method in a condition automatically adds "requires {http}" to the policy.

```
http-method [request* | proxy-request]
all STRING_COMPARISON
```

Example

```
http-method all values { POST }
```

http-version

Inspect the version of an HTTP request or response (since 11.4.0)

Specifying http-version in a condition automatically adds "requires {http}" to the policy.

```
http-version [proxy-request* | request | proxy-connect |
proxy-response | response]
[all STRING_COMPARISON]
[protocol STRING_COMPARISON]
[major NUMBER_COMPARISON]
[minor NUMBER_COMPARISON]
```

where

```
all - Compare against full version string
protocol - HTTP protocol
major - Numeric major part of HTTP version
minor - Numeric minor part of HTTP version
```

Examples

```
http-version all contains values { HTTP/1.1 }
http-version protocol values { HTTP }
http-version major values { 1 }
http-version minor greater-or-equal values { 0 }
```

http-status

Inspect the HTTP response's status (since 11.4.0)

Specifying http-status in a condition automatically adds "requires {http}" to the policy.

```
http-status [proxy-response* | response]
[all STRING_COMPARISON]
[code NUMBER_COMPARISON]
[text STRING_COMPARISON]
```

where

```
all - Compare against full HTTP status response
including both code and text
code - Numeric HTTP response status code
text - HTTP response status string, e.g.
"Authentication Required"
```

Examples

```
http-status response values { "HTTP/1.1 200 OK" }
http-status response code values { 401 }
http-status response text contains values { "Authentication Required" }
```

http-host

Match against an HTTP request's Host: header (since 11.4.0)

Specifying http-host in a condition automatically adds "requires {http}" to the policy.

```
http-host [proxy-request* | request | proxy-connect]
[all STRING_COMPARISON]
[host STRING_COMPARISON]
[port NUMBER_COMPARISON]
```

where

```
all - Compare against full Host header string
host - compare against hostname
port - compare against port number
```

Examples

```
http-host values { example.com }
http-host host values { ns1.example.com ns2.example.com }
http-host port values { 80 443 3128 8080 }
```

http-header

Match against any HTTP header (since 11.4.0)

Specifying http-header in a condition automatically adds "requires {http}" to the policy.

```
http-header [proxy-request* | request | proxy-connect |
proxy-response | response]
all STRING_COMPARISON
name STRING
```

Example

```
http-header response name Content-type starts-with values { text/ }
```

http-referer

Inspect HTTP Referer: header or specific parts of the URI value. (since 11.4.0)

Specifying http-referer in a condition automatically adds "requires {http}" to the policy.

```
http-referer [proxy-request* | request | proxy-connect]
[all STRING_COMPARISON
 [normalized BOOL]
]
[scheme STRING_COMPARISON
 [normalized BOOL]
]
[host STRING_COMPARISON
 [normalized BOOL]
]
[port NUMBER_COMPARISON
 [normalized BOOL]
]
[path STRING_COMPARISON
 [normalized BOOL]
]
[extension STRING_COMPARISON
 [normalized BOOL]
]
[query-string STRING_COMPARISON
 [normalized BOOL]
]
[path-segment STRING_COMPARISON
 index NUMBER
 [normalized BOOL]
]
[query-param STRING_COMPARISON
 name STRING
 [normalized BOOL]
]
[unnamed-query-param STRING_COMPARISON
 index NUMBER
 [normalized BOOL]
]
[urlcat
 [normalized BOOL]
]
```

where

all - entire URI, e.g.

http://example.com/path/to/page.cgi?a=b&c=d

normalized - Convert URI to standard form for consistent comparison.

scheme - e.g. http, https, ftp

normalized - Convert URI to standard form for consistent comparison.

host - DNS hostname or IP address

normalized - Convert URI to standard form for consistent comparison.

port - numeric port number, e.g. 80

normalized - Convert URI to standard form for consistent comparison.

path - URI path, e.g. /path/to

normalized - Convert URI to standard form for consistent comparison.

extension - document extension, e.g. cgi

normalized - Convert URI to standard form for consistent comparison.

query-string - full query string, e.g. a=b&c=d

normalized - Convert URI to standard form for consistent comparison.

path-segment - path segment by numerical index

index - Identify a segment of a path by its

numerical order starting at 1. Negative values

indicate counting right to left.

normalized - Convert URI to standard form for consistent comparison.

query-param - value of query param by name
 name - Identify a query string parameter by its name
 normalized - Convert URI to standard form for consistent comparison.
 unnamed-query-param - value of query parameter by numerical index
 index - Identify a query string parameter by its numerical order starting at 1. Negative values indicate counting right to left.
 normalized - Convert URI to standard form for consistent comparison.
 urlcat - Run URI through a categorization engine. List of categories - 'tmsh list sys url-db url-category'
 normalized - Convert URI to standard form for consistent comparison.

Examples

```

http-referer request all contains values { cgi }
http-referer request all scheme values { http https }
http-referer request all host values { example.com 127.0.0.1 }
http-referer request all port values { 80 8080 }
http-referer request all path contains values { /cgi-bin/ }
http-referer request all extension contains values { xml xhtml xsd }
http-referer request all query-string contains values { __utmz }
http-referer path-segment index 2 values { to }
http-referer query-parameter name foo contains values { bar }
http-referer request unnamed-query-param index 1 values { a }
  
```

http-cookie

Inspect an HTTP request's Cookie: header (since 11.4.0)

Specifying http-cookie in a condition automatically adds "requires {http}" to the policy.

```

http-cookie [proxy-request* | request | proxy-connect]
all STRING_COMPARISON
name STRING
  
```

Example

```
http-cookie name User values { xyz123 }
```

http-set-cookie

Inspect an HTTP response's Set-Cookie: header (since 11.4.0)

Specifying http-set-cookie in a condition automatically adds "requires {http}" to the policy.

```

http-set-cookie [proxy-response* | response]
[value STRING_COMPARISON
name STRING
]
[version STRING_COMPARISON
name STRING
]
[path STRING_COMPARISON
name STRING
]
[domain STRING_COMPARISON
name STRING
]
[expiry STRING_COMPARISON
name STRING
]
  
```

where

value - value of the named cookie named by the parameter
 version - version of the named cookie
 path - path of the named cookie
 domain - value of the domain specified by the named cookie
 expiry - Time when validity of named cookie expires, in RFC 6265 format (Wdy, DD Mon YYYY HH:MM:SS GMT)

Examples

```

http-set-cookie response value name Cust-Id values { org177 org187 org197 }
http-set-cookie response version name mycook values { 1.1 }
http-set-cookie response path name mycook values { /private/cgi-bin/ }
http-set-cookie response domain name mycook values { example.com }
http-set-cookie response expiry name MyCookie contains values { "Wed, 09 Jun 2021" }
  
```

http-basic-auth

Inspect an HTTP request's username/password specified for Basic authentication. (since 11.4.0)

Specifying http-basic-auth in a condition automatically adds "requires {http}" to the policy.

```
http-basic-auth [proxy-request* | request | proxy-connect]
[username STRING_COMPARISON]
[password STRING_COMPARISON]
```

where

username - basic authentication username
password - basic authentication password

Example

```
http-basic-auth password not values { password }
```

http-proxy

Inspect properties of the HTTP Explicit Proxy feature (since 13.1.0)

Specifying http-proxy in a condition automatically adds "requires {http-explicit}" to the policy.

```
http-proxy [request*]
[address IP_COMPARISON]
[port NUMBER_COMPARISON]
[route-domain NUMBER_COMPARISON]
```

where

address - The resolved IP address

Examples

```
http-proxy address equals values { 10.0.0.1 }
http-proxy port matches values { 80 }
http-proxy route-domain matches values { 2 }
```

http-connect

Inspect properties of the HTTP Proxy Connect feature (since 13.1.0)

Specifying http-connect in a condition automatically adds "requires {http-connect}" to the policy.

```
http-connect [client-accepted* | ssl-client-hello |
ssl-client-serverhello-send | ssl-server-hello |
ssl-server-handshake | server-connected |
proxy-request | request | proxy-connect |
proxy-response | response]
[host STRING_COMPARISON]
[port NUMBER_COMPARISON]
```

where

host - The host sent to the remote proxy

Examples

```
http-proxy-connect host matches values { http://example.com }
http-proxy-connect port matches values { 80 }
```

ssl-extension

Inspect SSL extensions being negotiated during HELLO phase. (since 11.4.0)

This condition is available to all policies.

```
ssl-extension [ssl-client-hello* | ssl-server-hello]
[server-name STRING_COMPARISON]
[npn STRING_COMPARISON
[index NUMBER]
]
[alpn STRING_COMPARISON
[index NUMBER]
]
```

where

server-name - server name indication
npn - next protocol negotiation
alpn - application layer protocol negotiation

Example

```
ssl-extension ssl-client-hello server-name values { secure43.example.org }
```

ssl-cert

Inspect properties of an SSL certificate. (since 11.4.0)

Specifying ssl-cert in a condition automatically adds "requires {server-ssl}" to the policy.

```
ssl-cert [ssl-server-handshake*]  
common-name STRING_COMPARISON  
[index NUMBER]
```

where

common-name - hostname covered by the SSL certificate

geoip

Specify a condition based upon properties of the geographical location of the IP address, such as continent code, country code, city, region, or organization. The default is to inspect the external interface, remote endpoint. (since 11.5.0)

This condition is available to all policies.

```
geoip [client-accepted* | ssl-client-hello |  
ssl-client-serverhello-send | ssl-server-hello |  
ssl-server-handshake | server-connected |  
proxy-request | request | proxy-connect |  
proxy-response | response]  
[continent STRING_COMPARISON  
[internal BOOL]  
[local BOOL]  
]  
[country-code STRING_COMPARISON  
[internal BOOL]  
[local BOOL]  
]  
[country-name STRING_COMPARISON  
[internal BOOL]  
[local BOOL]  
]  
[region-code STRING_COMPARISON  
[internal BOOL]  
[local BOOL]  
]  
[region-name STRING_COMPARISON  
[internal BOOL]  
[local BOOL]  
]  
[org STRING_COMPARISON  
[internal BOOL]  
[local BOOL]  
]  
[isp STRING_COMPARISON  
[internal BOOL]  
[local BOOL]  
]
```

where

continent - Two-character continent code: AF, AN, AS, OC, EU, NA, SA
country-code - Two-character country code as defined in ISO-3166-2
country-name - Full name of country
region-code - Abbreviation of State, Province, or country-specific region
region-name - Full name of State, Province, or country-specific region
org - Organization associated with address
isp - Internet Service Provider associated with address

Examples

```
geoip continent values { NA }  
geoip country-code values { us }  
geoip country-name values { "United States" }  
geoip region-code values { NY CA TX }  
geoip region-name values { Washington Oregon Idaho }  
geoip organization values { "Acme Widgets" }  
geoip isp values { "Fastcast Networks" "Responsive Cable Inc." }
```

cpu-usage

Specify a condition based upon CPU usage percentage for the past 15 seconds, 1 minute or 5 minutes intervals. (since 11.5.0)

This condition is available to all policies.

```
cpu-usage [client-accepted* | ssl-client-hello |
ssl-client-serverhello-send | ssl-server-hello |
ssl-server-handshake | server-connected |
proxy-request | request | proxy-connect |
proxy-response | response]
[last-15secs NUMBER_COMPARISON]
[last-1min NUMBER_COMPARISON]
[last-5mins NUMBER_COMPARISON]
```

where

last-15secs - CPU usage ratio % over the past 15 seconds, 0-100
last-1min - CPU usage ratio % over the past minute, 0-100
last-5mins - CPU usage ratio % over the past 5 minutes, 0-100

Examples

```
cpu-usage request last-15secs 8
cpu-usage response last-1min 10
cpu-usage last-5mins 12
```

http-user-agent

Specify a condition based upon User Agent sub-string, i.e. version, browser type, or mobile device make and model. (since 11.4.0)

Specifying http-user-agent in a condition automatically adds "requires {http}" to the policy.

```
http-user-agent [proxy-request* | request]
[device-make STRING_COMPARISON]
[device-model STRING_COMPARISON]
[browser-type STRING_COMPARISON]
[browser-version STRING_COMPARISON]
[ua-token STRING_COMPARISON]
name STRING
]
```

where

device-make - Make of device
device-model - Model of device
browser-type - Browser name/type
browser-version - Browser version string
ua-token - Sub version string associated with specified parameter

Examples

```
http-user-agent device-make values { Samsung ASUS }
http-user-agent device-model values { DroidX }
http-user-agent request browser-type values { Mozilla }
http-user-agent request browser-version values { "37.0.2049.0" }
http-user-agent user-agent-token name Mozilla values { 9.0 }
```

websocket

Specify a condition based upon properties of a websockets connection. (since 12.1.0)

Specifying websocket in a condition automatically adds "requires {websocket}" to the policy.

```
websocket [ws_request* | ws_response]
[protocol STRING_COMPARISON]
[extension STRING_COMPARISON]
[version STRING_COMPARISON]
[ws_key STRING_COMPARISON]
```

where

protocol - value of the Sec-WebSocket-Protocol header
extension - value of the Sec-WebSocket-Extensions header
version - value of the Sec-WebSocket-Version header
ws_key - value of the masking-key

classification

Specify a condition based on flow's classification results. (since 13.0.0)

Specifying classification in a condition automatically adds "requires {classification}" to the policy.

```
classification [classification-detected*]
[application STRING_COMPARISON]
[application-id ]
```

[application-risk]
[category STRING_COMPARISON]
[category-id]
[url-category STRING_COMPARISON]
[url-category-id]

where

application - Classification Application Name
application-id - Classification Application ID
application-risk - Classification Application Risk
category - Application's category name
category-id - Application's category ID
url-category - URL's category name
url-category-id - URL's category ID

Examples

application values { cnn youtube }
application-id values { 1245 }
application-risk greater-or-equal values { 3 }
category values { News_And_Media }
category-id values { 16666 }
url-category starts-with values { News }
url-category-id values { 25555 }

iprep

Perform a reputation lookup on IP address (since 13.1.0)

This condition is available to all policies.

```
iprep [client-accepted* | proxy-request | request |  
server-connected]  
all  
[internal BOOL]  
[local BOOL]
```

Example

```
iprep all values { "Spam Sources" "Denial of Service" }
```

ACTION_SPEC

An ACTION_SPEC, or action specification, is where you can tell the system the specific programmed actions you would like to take.

Actions are associated with an event, and depending on the action, can be set to run at different times during the request-response cycle.

Below is a list of all supported actions, the events during which they can be executed, their specific sub-actions, and parameters (if any).

Itm-policy

Provides the ability to disable LTM Policy processing on a request by request basis. (since 11.4.0)

This action is available to all policies.

```
Itm-policy [client-accepted* | ssl-client-hello |  
ssl-client-serverhello-send | ssl-server-hello |  
ssl-server-handshake | server-connected |  
proxy-request | request | proxy-connect |  
proxy-response | response]  
disable*
```

http

Provides the ability to enable or disable Big Ip's HTTP filter processing (since 11.4.0)

This action is available to all policies.

```
http [client-accepted* | proxy-request | request |  
response | server-connected]  
[enable*]  
[disable]
```

http-uri

Modify the request's URI or path or query string. Setting URI value overrides path and query-string setting. (since 11.4.0)

This action is available to all policies.

```
http-uri [proxy-request* | request | proxy-connect]  
replace*  
[value TCLSTRING]  
[path TCLSTRING]
```

[query-string TCLSTRING]

http-host

Modify the request's Host: header (since 11.4.0)

This action is available to all policies.

```
http-host [proxy-request* | request | proxy-connect]
  replace*
  value TCLSTRING
```

http-header

Modify HTTP header in request or response (since 11.4.0)

This action is available to all policies.

```
http-header [proxy-request* | request | proxy-connect |
proxy-response | response]
  [replace*
  name STRING
  value TCLSTRING
  ]
  [insert
  name STRING
  value TCLSTRING
  ]
  [remove
  name STRING
  ]
```

http-referer

Modify the request's Referer: header (since 11.4.0)

This action is available to all policies.

```
http-referer [proxy-request* | request | proxy-connect]
  [replace*
  [value TCLSTRING]
  ]
  [insert
  value TCLSTRING
  ]
  [remove]
```

where

replace - Replace request's Referer: header.
value - New value for request's Referer: header.
Values beginning with "tcl:" are treated as Tcl
command substitutions and expanded before use.

insert - Insert an HTTP Referer: header into the
request
value - New value for request's Referer: header.
Values beginning with "tcl:" are treated as Tcl
command substitutions and expanded before use.
remove - Remove the HTTP Referer: header from the
request

http-cookie

Modify the request's Cookie: header (since 11.4.0)

This action is available to all policies.

```
http-cookie [proxy-request* | request | proxy-connect]
  [insert*
  name STRING
  [value TCLSTRING]
  ]
  [remove
  name STRING
  ]
```

where

insert - Insert an HTTP Cookie: header into the
request
name - Name of the cookie being inserted into
request's Cookie: header.
value - New value for the cookie in the request's
Cookie: header. Values beginning with "tcl:" are
treated as Tcl command substitutions and expanded
before use.
remove - Remove the HTTP Cookie: header from the

request

name - Name of the cookie being inserted into request's Cookie: header.

Examples

```
http-cookie request insert name "Source-IP" value "tcl:[IP::remote_addr]"
http-cookie remove name "X-Tracker"
```

http-set-cookie

Modify the response's Set-Cookie: header (since 11.4.0)

This action is available to all policies.

```
http-set-cookie [response*]
[insert*
 name STRING
 value TCLSTRING
 [domain TCLSTRING]
 [path TCLSTRING]
]
[remove
 name STRING
]
```

where

insert - Insert an HTTP Set-Cookie: header into the response

name - Name of the cookie being inserted into response's Set-Cookie: header.
value - New value for the cookie in the response's Set-Cookie: header. Values beginning with "tcl:" are treated as Tcl command substitutions and expanded before use.

domain - Value for the domain attribute of a cookie in Set-Cookie header. Tcl command substitutions are allowed for this field.

path - Value for the path attribute of a cookie in Set-Cookie header. Tcl command substitutions are allowed for this field.

remove - Remove the HTTP Cookie: header from the request

request

name - Name of the cookie to be removed from response's Set-Cookie: header.

http-reply

Redirect an HTTP request to a different URL (since 11.4.0)

Specifying http-reply in an action automatically adds "controls {forwarding}" to the policy.

```
http-reply [proxy-request* | request | response]
redirect*
 location TCLSTRING
 [code NUMBER]
```

where

redirect - Redirect an HTTP request to a different URL

location - The new URL for which a redirect response will be sent. A Tcl command substitution can be used for this field.

code - Optional HTTP response code for redirect. Default value is 302 if not specified. Valid values are in the range of 300-399.

log

Write messages to local or remote system log (since 11.4.0)

This action is available to all policies.

```
log [client-accepted* | ssl-client-hello |
ssl-client-serverhello-send | ssl-server-hello |
ssl-server-handshake | server-connected |
proxy-request | request | proxy-connect |
proxy-response | response | ws_request |
ws_response | classification-detected]
write*
 [message TCLSTRING]
 [facility STRING]
 [priority STRING]
 [ip_address STRING]
 [port NUMBER]
```

where

write - Write a message to the system log files, local or remote

message - The message to write to the system log.

A Tcl command substitution is allowed here.

facility - Standard syslog facility associated with message, such as auth, kern, daemon, user.

priority - Standard syslog priority associated with message, such as emerg, alert, err, warning, notice, info, debug.

ip_address - For remote logging, the IP address of the remote syslog server.

port - For remote logging, the port number of the remote syslog server.

Examples

```
log request message "tcl:This is Tcl-enabled and the URI is [HTTP::uri]"
```

```
log request message "This is just a plain old string"
```

```
log request message "Something serious is happening!" priority alert
```

pem

Classify traffic (since 11.4.0)

Specifying pem in an action automatically adds "controls {classification}" to the policy.

```
pem [ssl-client-hello* | ssl-client-serverhello-send |  
ssl-server-hello | ssl-server-handshake |  
proxy-request | request | response]  
classify*  
  [application STRING]  
  [category STRING]  
  [protocol STRING]  
  [defer BOOL]  
  [ssl-session-id BOOL]
```

where

classify - Classify traffic

ce

Classify traffic (since 12.1.0)

Specifying ce in an action automatically adds "controls {ce}" to the policy.

```
ce [ssl-client-hello* | ssl-server-hello |  
ssl-server-handshake | request | response]  
classify*  
  [application STRING]  
  [category STRING]
```

where

classify - Classify traffic

cache

Control caching (since 11.4.0)

Specifying cache in an action automatically adds "controls {caching}" to the policy.

```
cache [client-accepted* | proxy-request | request |  
response | server-connected]  
[enable*  
  [pin BOOL]  
]  
[disable]
```

where

enable - Enable caching for a connection

disable - Disable caching for a connection

compress

Control compression (since 11.4.0)

Specifying compress in an action automatically adds "controls {compression}" to the policy.

```
compress [client-accepted* | proxy-request | request |  
response | server-connected]  
[enable*]  
[disable]
```

where

enable - Enable compression for a connection
disable - Disable compression for a connection

decompress

Control decompression (since 11.4.0)

Specifying decompress in an action automatically adds "controls {compression}" to the policy.

```
decompress [client-accepted* | proxy-request | request |  
response | server-connected]  
[enable*]  
[disable]
```

where

enable - Enable decompression for a connection
disable - Disable decompression for a connection

forward

Many options for controlling where a connection is forwarded (since 11.4.0)

Specifying forward in an action automatically adds "controls {forwarding}" to the policy.

```
forward [client-accepted* | ssl-client-hello |  
ssl-client-serverhello-send | proxy-request |  
request]  
[select*  
[pool STRING]  
[fallback-pool STRING]  
[clone-pool STRING]  
[node STRING]  
[snat STRING]  
[snatpool STRING]  
[nexthop STRING]  
[vlan STRING]  
[vlan-id NUMBER]  
[virtual STRING]  
[rateclass STRING]  
]  
[reset]
```

where

select - Select appropriate location for forwarding the connection based on specified parameters. While all of the parameters are marked as optional, at least one must be specified.

pool - Forward connection to the specified pool.

fallback-pool - Forward connection to the specified pool when the default pool does not have active members.

clone-pool - Clone traffic to the specified clone pool.

node - Forward connection to the specified node.

snat - Control snat automap.

snatpool - Forward connection to the specified snat pool.

nexthop - Set the next destination for the connection to the specified endpoint. A vlan or vlan-id must also be specified.

vlan - Forward connection to the specified vlan.

vlan-id - Forward connection to the vlan specified by the vlan-id.

virtual - Forward connection to the specified virtual server.

rateclass - Control rate class properties on the connection.

reset - Deprecated. See target "shutdown connection".

shutdown

Reset connection (since 13.1.0)

This action is available to all policies.

```
shutdown [client-accepted* | ssl-client-hello |  
ssl-client-serverhello-send | ssl-server-hello |  
ssl-server-handshake | server-connected |  
proxy-request | request | proxy-connect |  
proxy-response | response | ws_request |  
ws_response]  
connection*
```

where

connection - Terminate the connection through the Big IP.

persist

Many options for controlling how a connection is persisted (since 12.0.0)

Specifying persist in an action automatically adds "controls {persistence}" to the policy.

```
persist [client-accepted* | proxy-request | request]
[disable]
[src-addr*
 [netmask STRING]
 [timeout NUMBER]
 ]
[dest-addr
 [netmask STRING]
 [timeout NUMBER]
 ]
[cookie-insert
 [name STRING]
 [expiry STRING]
 ]
[cookie-rewrite
 [name STRING]
 [expiry STRING]
 ]
[cookie-passive
 [name STRING]
 ]
[cookie-hash
 name STRING
 [offset NUMBER]
 [length NUMBER]
 [timeout NUMBER]
 ]
[universal
 key TCLSTRING
 [timeout NUMBER]
 ]
[hash
 key TCLSTRING
 [timeout NUMBER]
 ]
[carp
 key TCLSTRING
 [timeout NUMBER]
 ]
```

where

disable - Disable persistence

src-addr - Persist the connection based on the source IP address.

netmask - Network mask, e.g. 192.168.13.23/16 or 10.0.2.15/255.0.0.0.

timeout - Timeout value in seconds.

dest-addr - Persist the connection based on the destination IP address.

netmask - Network mask, e.g. 192.168.13.23/16 or 10.0.2.15/255.0.0.0.

timeout - Timeout value in seconds.

cookie-insert - Persist the connection using cookie insertion method

name - cookie name

expiry - Expiration duration expressed as [Dd]

[[HH:]MM:]SS

cookie-rewrite - Persist the connection using cookie rewrite method

name - cookie name

expiry - Expiration duration expressed as [Dd]

[[HH:]MM:]SS

cookie-passive - Persist the connection using cookie passive method

name - cookie name

cookie-hash - Persist the connection using cookie hash method

name - cookie name

offset - offset into hash

length - substring length

timeout - Timeout value in seconds.

universal - persistence based on user-defined key

key - The key to use. Tcl command substitution is allowed.

timeout - Timeout value in seconds.

hash - persistence based on hash of the key
key - The key to use. Tcl command substitution is allowed.
timeout - Timeout value in seconds.
carp - hash persistence using Cache Array Routing Protocol (CARP) algorithm
key - The key to use. Tcl command substitution is allowed.
timeout - Timeout value in seconds.

wam

Control web acceleration (since 11.4.0)

Specifying wam in an action automatically adds "controls {wam}" to the policy.

```
wam [client-accepted* | proxy-request | request]
[enable*]
[disable]
```

where

enable - Enable web acceleration for a connection
disable - Disable web acceleration for a connection

asm

Control web security (since 11.4.0)

Specifying asm in an action automatically adds "controls {asm}" to the policy.

```
asm [client-accepted* | proxy-request | request]
[enable*]
policy STRING
]
[disable]
```

where

enable - Enable web security for a connection
policy - name of security policy to enable
disable - disable web security for the connection

l7dos

Enable or disable Layer 7 Denial-of-Service processing (since 11.4.0)

Specifying l7dos in an action automatically adds "controls {l7dos}" to the policy.

```
l7dos [client-accepted* | proxy-request | request]
[enable*]
[from-profile STRING]
]
[disable]
```

where

enable - turn on Layer 7 DOS protection
from-profile - name of DOS profile to enable
disable - turn off Layer 7 DOS protection

bot-defense

Enable or disable Unified Bot Defense processing (since 14.1.0)

Specifying bot-defense in an action automatically adds "controls {bot-defense}" to the policy.

```
bot-defense [client-accepted* | proxy-request | request]
[enable*]
[from-profile STRING]
]
[disable]
```

where

enable - turn on Bot Defense protection
from-profile - name of bot-defense profile to enable
enable
disable - turn off Bot Defense protection

avr

Enable or disable Application Visibility and Reporting (since 11.4.0)

Specifying avr in an action automatically adds "controls {avr}" to the policy.

```
avr [client-accepted* | proxy-request | request]
```

[enable*]
[disable]

where

enable - turn on reporting
disable - turn off reporting

tcl

Set a Tcl variable in runtime environment (since 11.4.0)

This action is available to all policies.

```
tcl [client-accepted* | ssl-client-hello |  
ssl-client-serverhello-send | ssl-server-hello |  
ssl-server-handshake | server-connected |  
proxy-request | request | proxy-connect |  
proxy-response | response]  
  set-variable*  
    name STRING  
    expression STRING
```

where

set-variable - set a Tcl variable in the runtime
environment
name - name of variable
expression - Tcl expression to evaluate

Example

```
tcl set-variable expression tcl:[HTTP::uri] name my_uri
```

request-adapt

Enable or disable request adaptation, optionally sending traffic to specified internal virtual server (since 11.4.0)

Specifying request-adapt in an action automatically adds "controls {request-adaptation}" to the policy.

```
request-adapt [client-accepted* | server-connected |  
proxy-request | request | response]  
  [enable*  
  [internal-virtual STRING]  
  ]  
  [disable]
```

where

enable - turn on request adaptation
internal-virtual - which internal virtual server
disable - turn off request adaptation

response-adapt

Enable or disable response adaptation, optionally sending traffic to specified internal virtual server (since 11.4.0)

Specifying response-adapt in an action automatically adds "controls {response-adaptation}" to the policy.

```
response-adapt [client-accepted* | server-connected |  
proxy-request | request | response]  
  [enable*  
  [internal-virtual STRING]  
  ]  
  [disable]
```

where

enable - turn on response adaptation
internal-virtual - which internal virtual server
disable - turn off response adaptation

tcp-nagle

Enable or disable Nagle's algorithm on a connection, or allow BIG-IP to determine the best setting. (since 11.4.0)

This action is available to all policies.

```
tcp-nagle [client-accepted* | request]  
  [enable*  
  [auto-mode BOOL]  
  ]  
  [disable]
```

where

enable - turn on Nagle
auto-mode - turn on Auto Nagle
disable - turn off Nagle

server-ssl

Enable or disable encrypted connections to backend servers (since 11.4.0)

Specifying server-ssl in an action automatically adds "controls {server-ssl}" to the policy.

server-ssl [client-accepted* | proxy-request | request |
proxy-connect | proxy-response | server-connected]
[enable*]
[disable]

where

enable - encrypted connection to backend
disable - plaintext connection to backend

client-ssl

Enable or disable encrypted connections to client (since 13.1.0)

Specifying client-ssl in an action automatically adds "controls {client-ssl}" to the policy.

client-ssl [client-accepted*]
[enable*]
[disable]

where

enable - encrypted connection to client
disable - plaintext connection to client

ssl-intercept

Switch the SSL Intercept feature between bypass and intercept modes (since 13.1.0)

Specifying ssl-intercept in an action automatically adds "controls {ssl-intercept}" to the policy.

ssl-intercept [ssl-client-serverhello-send*]
[enable*]
[disable]

where

enable - SSL Intercept feature is in intercept mode
disable - SSL Intercept feature is in bypass mode

websocket

Enable or disable websocket processing (since 12.1.0)

Specifying websocket in an action automatically adds "controls {websocket}" to the policy.

websocket [client-accepted* | ws_request | ws_response |
proxy-request | request | proxy-connect |
proxy-response | response | server-connected]
[enable*]
[disable]

where

enable - turn on websocket filter
disable - turn off websocket filter

classification

Perform classification/enforcement of the connection (since 13.0.0)

Specifying classification in an action automatically adds "controls {classification}" to the policy.

classification [ssl-client-hello* | ssl-client-serverhello-send |
ssl-server-hello | ssl-server-handshake |
proxy-request | request | response | classification-detected]
[drop]
[reject]

where

drop - Silently drop connection
reject - Gracefully close connection

server-ssl-profile

Dynamic selection of SSL server profile (since 13.1.0)

Specifying server-ssl-profile in an action automatically adds "controls {server-ssl}" to the policy.

```
server-ssl-profile [server-connected*]
select*
name STRING
```

http-proxy

Enable or Disable the HTTP Explicit Proxy (since 13.1.0)

This action is available to all policies.

```
http-proxy [proxy-request*]
[enable*]
[disable]
```

where

enable - turn on the HTTP Explicit Proxy feature
disable - turn off the HTTP Explicit Proxy feature

http-uri-rewrite

Enable or Disable the rewriting of HTTP URI's into proxy form (since 13.1.0)

This action is available to all policies.

```
http-uri-rewrite [client-accepted* | ssl-client-hello |
ssl-server-hello | ssl-server-handshake |
proxy-request | request | server-connected]
[enable*]
[disable]
```

where

enable - turn on rewriting the URI into proxy form
disable - turn off rewriting the URI into proxy form

http-connect

Control the HTTP Proxy Connect feature (since 13.1.0)

This action is available to all policies.

```
http-connect [client-accepted* | ssl-client-hello |
proxy-request | request | server-connected |
proxy-connect | proxy-response]
[enable*]
[disable]
[replace
[port NUMBER]
[host STRING]
]
[retry]
```

where

enable - turn on the HTTP Proxy Connect feature
disable - turn off the HTTP Proxy Connect feature
retry - retry the HTTP CONNECT

SEE ALSO

create, delete, edit, glob, list, modify, ltm policy-strategy, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2016-2020. All rights reserved.

BIG-IP 2020-06-23 ltm policy(1)

ltm pool

NAME

pool - Configures load balancing pools for the Local Traffic Manager.

MODULE

ltm

SYNTAX

Modify the pool component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create pool [name]
modify pool [name]
options:
  all
  allow-nat [yes | no]
  allow-snat [yes | no]
  app-service [[string] | none]
  autoscale-group-id [[string] | none]
  description [string]
  gateway-failsafe-device [string]
  ignore-persisted-weight [yes | no]
  ip-tos-to-client [pass-through | [integer] ]
  ip-tos-to-server [pass-through | [integer] ]
  link-qos-to-client [pass-through | [integer] ]
  link-qos-to-server [pass-through | [integer] ]
  load-balancing-mode [dynamic-ratio-member | dynamic-ratio-node |
    fastest-app-response | fastest-node |
    least-connections-members |
    least-connections-node |
    least-sessions |
    observed-member | observed-node |
    predictive-member | predictive-node |
    ratio-least-connections-member |
    ratio-least-connections-node |
    ratio-member | ratio-node | ratio-session |
    round-robin | weighted-least-connections-member |
    weighted-least-connections-node]
  members [add | delete | modify | replace-all-with] {
    [ [node_name:port] ] {
      options:
        address [ip address]
        app-service [[string] | none]
        connection-limit [integer]
        description [string]
        dynamic-ratio [integer]
        inherit-profile [enabled | disabled]
        logging [enabled | disabled]
        monitor [name]
        priority-group [integer]
        profiles [none | profile_name]
        rate-limit [integer]
        ratio [integer]
        session [user-enabled | user-disabled]
        state [ user-up | user-down ]
        fqdn {
          name [string]
          autopopulate [enabled | disabled]
        }
      }
    }
  }
  members none
  metadata
    [add | delete | modify] {
      [metadata_name ... ] {
        value [ "value content" ]
      }
    }
  persist [ true | false ]
}
min-active-members [integer]
min-up-members [integer]
min-up-members-action [failover | reboot | restart-all]
min-up-members-checking [enabled | disabled]
monitor [name]
profiles [none | profile_name]
queue-on-connection-limit [enabled | disabled]
queue-depth-limit [integer]
queue-time-limit [integer]
reselect-tries [integer]
service-down-action [drop | none | reselect | reset]
slow-ramp-time [integer]

edit pool [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

reset-stats pool
reset-stats pool [ [ [name] | [glob] | [regex] ] ... ]
```

```

mv pool [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
options:
  to-folder

DISPLAY
list pool
list pool [ [ [name] | [glob] | [regex] ] ... ]
show running-config pool
show running-config pool [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

show pool
show pool [name]
options:
  all-properties
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  detail
  field-fmt

DELETE
delete pool [name]

```

Note: You must remove all references to a pool before you can delete the pool.

DESCRIPTION

You can use this pool component to configure the pool definitions on the Local Traffic Manager. A load balancing pool is a logical set of devices, such as Web servers, that you group together to receive and process traffic.

EXAMPLES

```
create pool my_pool members add { 10.2.3.11:http 10.2.3.12:http }
```

Creates a Local Traffic Manager pool named my_pool with two members, 10.2.3.11 and 10.2.3.12, using the default values for the pool and pool members.

```
delete pool my_pool
```

Deletes the pool named my_pool.

```
show pool
```

Displays statistics and status for all Local Traffic Manager pools in the system configuration.

```
show pool all-properties
```

Displays statistics and status for all Local Traffic Manager pools in the system configuration.

Note that if the system includes Packet Velocity(r) ASIC (PVA) and PVA Assist capabilities, this command displays status and statistics for that feature.

```
list pool my_pool
```

Displays properties of the pool named my_pool.

```
mv /ltm pool /Common/my_pool to-folder /Common/some_folder
```

Moves an LTM pool named my_pool and all of its Pool Members to the folder named some_folder, where some_folder has already been created under /Common.

Please refer to the mv manual page for additional examples on how to use the mv command.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

allow-nat

Specifies whether the pool can load balance network address translation (NAT) connections. The default value is yes.

allow-snat

Specifies whether the pool can load balance secure network address translation (SNAT) connections. The default value is yes.

app-service

Specifies the name of the application service to which the pool belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the pool. Only the application service can modify or delete the pool.

autoscale-group-id

Deprecated since v13.1.0. Specifies the autoscale-group id as reported by Amazon Web Services(AWS).

description

User defined description.

gateway-failsafe-device

Specifies that the pool is a gateway failsafe pool in a redundant configuration. The gateway-failsafe-device identifies the device that depends on the gateway. If the monitor associated with the pool reports that the gateway is down, the device goes to the standby state. The default value for this string is empty, the feature is not configured.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

traffic-acceleration-status

Displays the current Traffic-Acceleration status. Indicates whether the pool is in-use by or dedicated to a virtual server that uses a traffic-acceleration profile.

ignore-persisted-weight

Discounts the weight of connections made to pool members selected through persistence, rather than as a result of the algorithm configured on the pool. If the connection's weight is ignored, then it is not treated as a 'pick' for that pool member, and does not influence subsequent pool member load balancing decisions.

This option only impacts pools configured with one of the following load balancing modes: observed-member, observed-node, predictive-member, predictive-node, ratio-least-connections-member, ratio-least-connections-node, ratio-member, or ratio-node.

The default value is no, which results in persisted pool member connections being accounted for during load balancing calculations.

ip-tos-to-client

Specifies the Type of Service (ToS) level to use when sending packets to a client. The default value is 65535 (pass-through).

ip-tos-to-server

Specifies the ToS level to use when sending packets to a server. The default value is 65535 (pass-through).

link-qos-to-client

Specifies the Link Quality of Service (QoS) level to use when sending packets to a client. The default value is 65535 (pass-through).

link-qos-to-server

Specifies the Link QoS level to use when sending packets to a server. The default value is 65535 (pass-through).

load-balancing-mode

Specifies the modes that the system uses to load balance name resolution requests among the members of this pool. The default value is round-robin.

The options are:

dynamic-ratio-member

Specifies that the system distributes connections based on various aspects of real-time server performance analysis, such as the number of current connections per node or the fastest node response time.

This mode is similar to the dynamic-ratio-node mode, except that weights are based on continuous monitoring of the servers and are therefore continually changing.

dynamic-ratio-node

Specifies that the system distributes connections based on various aspects of real-time server performance analysis, such as the number of current connections per node or the fastest node response time.

This mode is similar to the dynamic-ratio-member mode, except that weights are based on continuous monitoring of the servers and are therefore continually changing.

fastest-app-response

Specifies that the system passes a new connection based on the fastest response of all currently active nodes in a pool. This mode might be particularly useful in environments where nodes are distributed across different logical networks.

fastest-node

Specifies that the system passes a new connection based on the fastest response of all pools of which a server is a member. This mode might be particularly useful in environments where nodes are distributed across different logical networks.

least-connections-member

Specifies that the system passes a new connection to the node that has the least number of current connections in the pool. This mode works best in environments where the servers or other equipment you are load balancing have similar capabilities.

This dynamic load balancing mode distributes connections based on various aspects of real-time server performance analysis, such as the current number of connections per node or the fastest node response time.

least-connections-node

Specifies that the system passes a new connection to the node that has the least number of current connections out of all pools of which a node is a member. This mode works best in environments where

the servers or other equipment you are load balancing have similar capabilities.

This dynamic load balancing mode distributes connections based on various aspects of real-time server performance analysis, such as the number of current connections per node, or the fastest node response time.

least-sessions

Specifies that the system passes a new connection to the node that has the least number of current sessions. This mode works best in environments where the servers or other equipment you are load balancing have similar capabilities.

This dynamic load balancing mode distributes connections based on various aspects of real-time server performance analysis, such as the number of current sessions.

observed-member

Specifies that the system ranks nodes based on the number of connections. Nodes that have a better balance of fewest connections receive a greater proportion of the connections.

This mode differs from the least-connections-member mode, which measures connections only at the moment of load balancing, while the observed-member mode tracks the number of Layer 4 connections to each node over time and creates a ratio for load balancing.

This dynamic load balancing mode works well in any environment, but may be particularly useful in environments where node performance varies significantly.

observed-node

Specifies that the system ranks nodes based on the number of connections. Nodes that have a better balance of fewest connections receive a greater proportion of the connections.

This mode differs from least-connections-node mode, which measures connections only at the moment of load balancing, while the observed-node mode tracks the number of Layer 4 connections to each node over time and creates a ratio for load balancing.

This dynamic load balancing method works well in any environment, but may be particularly useful in environments where node performance varies significantly.

predictive-member

Uses the ranking method used by the observed-member mode, except that the system analyzes the trend of the ranking over time, determining whether a node's performance is improving or declining. The nodes in the pool with better performance rankings that are currently improving, rather than declining, receive a higher proportion of the connections. This dynamic load balancing mode works well in any environment.

predictive-node

Uses the ranking method used by the observed-node mode, except that the system analyzes the trend of the ranking over time, determining whether a node's performance is improving or declining. The nodes in the pool with better performance rankings that are currently improving, rather than declining, receive a higher proportion of the connections. This dynamic load balancing mode works well in any environment.

queue-depth-limit

Specifies the maximum number of connections that may simultaneously be queued to go to any member of this pool. The default is zero which indicates there is no limit.

queue-on-connection-limit

Enable or disable queuing connections when pool member or node connection limits are reached. When queuing is not enabled, new connections are reset when connection limits are met. The default value is disabled.

queue-time-limit

Specifies the maximum time, in milliseconds, a connection will remain enqueued. The default is zero which indicates there is no limit.

ratio-least-connections-member

Specifies that the system weights connections to each pool member based on the value of the ratio weight defined for each pool member. If a ratio weight is unspecified, it will be treated as a default value of '1'.

ratio-least-connections-node

Specifies that the system weights connections to each pool member based on the value of the ratio weight defined for the pool member's node. If a ratio weight is unspecified, it will be treated as a default value of '1'.

ratio-member

Specifies that the number of connections that each machine receives over time is proportionate to a ratio weight you define for each machine within the pool.

ratio-node

Specifies that the number of connections that each machine receives over time is proportionate to a ratio weight you define for each machine across all pools of which the server is a member.

ratio-session

Specifies that the number of sessions that each machine receives over time is proportionate to a ratio weight that you define for each machine within the pool.

round-robin

Specifies that the system passes each new connection request to the next server in line, eventually

distributing connections evenly across the array of machines being load balanced. This mode works well in most configurations, especially if the equipment that you are load balancing is roughly equal in processing speed and memory.

weighted-least-connections-member

Specifies that the system passes a new connection to the pool member that is handling the lowest percentage of the specified maximum number of concurrent connections allowed. This mode works best in environments where the servers or other equipment you are load balancing have different but quantified capability limits.

This mode requires that you specify a value for the connection-limit option for all members of the pool, but does not require all servers or other equipment you are load balancing to have similar capabilities.

weighted-least-connections-node

Specifies that the system passes a new connection to the node that is handling the lowest percentage of the specified connection limit. This mode works best in environments where the servers or other equipment you are load balancing have different but quantified capability limits.

This mode requires that you specify a value for the connection-limit option for all nodes, but does not require all servers or other equipment you are load balancing to have similar capabilities.

members

Adds, deletes, or replaces a set of pool members, by specifying a node name and service port in the format [node name/port]. If a node by the specified name does not exist, it will be created. You can configure the following options for a pool member:

address

Specifies the IP address of a pool member if a node by the name specified does not already exist.

app-service

Specifies the name of the application service to which the pool member belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the pool member. Only the application service can modify or delete the pool member.

connection-limit

Specifies the maximum number of concurrent connections allowed for a pool member. The default value is 0 (zero).

description

User defined description.

dynamic-ratio

Specifies a range of numbers that you want the system to use in conjunction with the ratio load balancing method. The default value is 1.

fqdn Specifies the attributes for defining a fully qualified domain name for the node.

name Specifies the fully-qualified domain name of the node.

address-family

Specifies whether the fqdn should consider IPv4, IPv6, or IP-agnostic address family.

autopopulate

Specifies whether a node defined by a fully-qualified domain name should automatically scale to the set of IP addresses returned by the DNS query. If disabled, only one ephemeral node is generated from the first IP address returned by DNS. The default is disabled.

interval

Specifies the interval, in seconds, to instantiate DNS queries on a fully-qualified domain name. The default is 3600. A value of 'ttl' uses the TTL value obtained from the DNS server.

down-interval

Specifies the interval for the domain name resolution operation when a DNS query fails. The default is 5.

inherit-profile

Specifies whether the pool member inherits the encapsulation profile from the parent pool. The default value is enabled. If you disable inheritance, no encapsulation takes place, unless you specify another encapsulation profile for the pool member using the profiles attribute.

logging

Specifies whether the monitor applied should log its actions. Logs are stored in /var/log/monitors/ and are regularly rotated and compressed. The default value is disabled. This option isn't a part of configuration and will reset to disabled on load. This option doesn't sync.

monitor

Specifies the health monitors that are configured to monitor the pool member. The default value is default, the system monitors the pool member using the monitors specified for the pool.

You can specify:

• A single monitor, for example, modify pool mypool members modify { pool_member_1:80 { monitor http } }.

• Multiple monitors, for example, modify pool mypool members modify { pool_member_1:80 { monitor

http and https } }.

Â A minimum number of monitors, for example, modify pool mypool members modify { pool_member_1:80 { monitor min 1 of { http https } } }.

Â No monitor rule or remove a monitor rule, for example, modify pool mypool members modify { pool_member_1:80 { monitor none } }.

profiles

Specifies the encapsulation profile to use for the pool member, when the inherit-profile attribute is disabled. The default value is none.

priority-group

Specifies the priority group within the pool for this pool member. Valid values are 0 through 65535. The system sends traffic to groups in order of priority. The default value is 0.

rate-limit

Specifies the maximum number of connections per second allowed for a pool member. The default value is 'disabled'.

ratio

Specifies the weight of the pool member for load balancing purposes. The default value is 1.

session

Establishing a session with a pool member is establishing the ability of the client to persist to the pool member when making new connections. When a pool member is session disabled, clients that have already established sessions with the pool member may create new connections, but a client that has not already established a session may not create a new one (or make a connection which would create a new session). This feature is used to gently drain connections from a node, typically as part of a maintenance operation. The default value is user-enabled.

The value of this property can be set by system or by user. If the value is set by system, the property will not be displayed in "Edit" command. But, users can add this field in if they need to modify this property. The values which user can set for this property are user-enabled and user-disabled.

state

user-down forces the pool member offline, overriding monitors. user-up reverts the user-down. When user-up, this displays the monitor state.

metadata

Associates user-defined data, each of which has name and value pair and persistence. The default value is persistent, which saves the data to the config file.

min-active-members

Specifies the minimum number of members that must be up for traffic to be confined to a priority group when using priority-based activation. The default value is 0 (zero). An active member is a member that is up (not marked down) and is handling fewer connections than its connection limit.

min-up-members

Specifies the minimum number of pool members that must be up; otherwise, the system takes the action specified in the min-up-members-action option.

Use this option for gateway pools in a redundant system where a unit number is applied to the pool. This indicates that the pool is configured only on the specified unit.

min-up-members-action

Specifies the action to take if min-up-members-checking is enabled, and the number of active pool members falls below the number specified in the min-up-members option. The default value is failover. The options are:

reboot

Specifies that when the min-up-members-checking option is enabled, and the number of active pool members is less than the number specified in the min-up-members option, the system restarts.

restart-all

Specifies that when the min-up-members-checking option is enabled, and the number of active pool members is less than the number specified in the min-up-members option, the system restarts.

failover

Specifies, for a redundant system, that when the min-up-members-checking option is enabled, and the number of active pool members is less than the number specified in the min-up-members option, the system fails over.

min-up-members-checking

Enables or disables the min-up-members feature. If you enable this feature, you must also specify a value for both the min-up-members and min-up-members-action options.

monitor

Specifies the health monitors that the system uses to determine whether it can use this pool for load balancing. The monitor marks the pool up or down based on whether the monitor is successful. The default value is none.

You can specify:

Â A single monitor, for example, modify pool mypool monitor http.

Â· Multiple monitors, for example, modify pool mypool monitor http and https.

Â· A minimum number of monitors, for example, modify pool mypool monitor min 1 of {http and https}.

Â· No monitor rule or remove a monitor rule, for example, modify pool mypool monitor none.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the pool resides.

profiles
Specifies the profile to use for encapsulation. The default value is none, which indicates no encapsulation.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reselect-tries
When set to the default value of 0 (zero) the system does not attempt to load balance to another pool member after a passive failure. A passive failure is a pool member connection failure.

When set to any other value, the system attempts to load balance to another pool member after a passive failure, and if that attempt also results in a passive failure, the system repeats the process until the specified number of reselection tries is reached.

reset-stats
Resets the statistics for the specified component to 0 (zero).

service-down-action
Specifies the action to take if the service specified in the pool is marked down. The options are:

drop Specifies that the system drops connections when a the service is marked down.

none Specifies that the system takes no action when a the service is marked down. This is the default value.

reselect
Specifies that the system reselects a node for the next packet that comes in on a Layer 4 connection, if the service of the existing connection is marked down.

reset
Specifies that the system resets when a the service is marked down.

slow-ramp-time
Specifies, in seconds, the ramp time for the pool. This provides the ability to cause a pool member that has just been enabled, or marked up, to receive proportionally less traffic than other members in the pool. The proportion of traffic the member accepts is determined by how long the member has been up in comparison to the value of the slow-ramp-time option for the pool.

For example, if the load-balancing-mode of a pool is round-robin and it has a slow-ramp-time of 60 seconds, when a pool member has been up for only 30 seconds, the pool member receives approximately half the amount of new traffic as other pool members that have been up for more than 60 seconds. After the pool member has been up for 45 seconds, it receives approximately three quarters of the new traffic.

The slow-ramp-time option is particularly useful when used with the least-connections-member load balancing mode. The default value is 10.

to-folder
This is used with the mv command to specify a folder in which to move the pool and its members to.

Note: pools can be moved to any folder under /Common, but dependencies upon it may restrict moving it out of /Common.

SEE ALSO
create, delete, edit, glob, list, modify, mv, ltm virtual, regex, reset-stats, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

BIG-IP 2018-10-17 ltm pool(1)

NAME

analytics - Configures an analytics profile.

MODULE

lrm profile

SYNTAX

Configure the analytics component within the lrm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create analytics [name]

modify analytics [name]

options:

alerts [none | add | delete | modify | replace-all-with] {
 name [string] {

options:

app-service [[string] | none]

granularity [application | pool-member |
 virtual-server]

metric [average-page-load-time | average-request-throughput |
 average-response-throughput | average-server-latency |
 average-tps | max-page-load-time | max-request-throughput |
 max-server-latency | max-response-throughput | max-tps]

sample-period [integer]

threshold [integer]

threshold-relation [above | below]

}

}

app-service [[string] | none]

captured-traffic-external-logging [enabled | disabled]

captured-traffic-internal-logging [enabled | disabled]

collect-page-load-time [enabled | disabled]

collect-geo [enabled | disabled]

collect-http-throughput [enabled | disabled]

collect-http-timing-metrics [enabled | disabled]

collect-ip [enabled | disabled]

collect-max-tps-and-throughput [enabled | disabled]

collect-methods [enabled | disabled]

collect-response-codes [enabled | disabled]

collect-server-latency [enabled | disabled]

collect-subnets [enabled | disabled]

collect-url [enabled | disabled]

collect-user-agent [enabled | disabled]

collect-user-sessions [enabled | disabled]

collected-stats-external-logging [enabled | disabled]

collected-stats-internal-logging [enabled | disabled]

defaults-from [analytics profile name [string] | none]

description [string]

external-logging-publisher [name]

notification-by-email [enabled | disabled]

notification-by-snmp [enabled | disabled]

notification-by-syslog [enabled | disabled]

notification-email-addresses [none | add | delete | modify |

 replace-all-with] { email-address [string] }

publish-irule-statistics [enabled | disabled]

sampling [enabled | disabled]

session-cookie-security [always-secure | ssl-only | never-secure]

session-timeout-minutes [integer]

smtp-config [smtp configuration object name]

subnet-masks [none | add | delete | modify |

 replace-all-with] {

 name [string] {

options:

 subnet [IPv4/IPv6 address]

}

countries-for-stat-collection [add | delete]

ips-for-stat-collection [add | delete]

subnets-for-stat-collection [add | delete]

urls-for-stat-collection [add | delete]

}

traffic-capture [none | add | delete | modify |

 replace-all-with] {

 name [string] {

options:

 app-service [[string] | none]

 captured-protocols [all | http | https]

 client-ips [none | add | delete | modify |

 replace-all-with] { ipv4.address }

 dos-activity [any | mitigated-by-dos17]

 methods [none | add | delete | modify |

 replace-all-with] { method [string] }

 node-addresses [none | add | delete | modify |

 replace-all-with] { node }

 request-captured-parts [all | body | headers | none]

 request-content-filter-search-part [all | body | headers |

 none | uri]

```
request-content-filter-search-string [none | [string]]
response-captured-parts [all | body | headers | none]
response-codes [none | add | delete | modify |
  replace-all-with] { response-code [integer] }
response-content-filter-search-part [all | body |
  headers | none]
response-content-filter-search-string [none | [string]]
url-filter-type [all | black-list | white-list]
url-path-prefixes [none | add | delete | modify |
  replace-all-with] { url-path-prefix [string] }
user-agent-substrings [none | add | delete | modify |
  replace-all-with] { user-agent-substring [string] }
virtual-servers [none | add | delete | modify |
  replace-all-with] { virtual }
}
```

```
edit analytics [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list analytics
list analytics [ [ [name] | [glob] | [regex] ] ... ]
show running-config analytics
show running-config analytics [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete analytics [name]
```

DESCRIPTION

Use the analytics component to create, modify, display, or delete an analytics profile for use with analytics functionality.

EXAMPLES

```
create analytics my_analytics_profile defaults-from analytics
```

Creates a custom analytics profile named my_analytics_profile that inherits its settings from the system default analytics profile.

```
list analytics
```

Displays the properties of all analytics profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

alerts

Adds, deletes, or replaces a set of analytics alerts. You can configure the following options for an analytics alert:

app-service

Specifies the name of the application service to which the alert belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the alert. Only the application service can modify or delete the alert.

granularity

Specifies a granularity level on which the alert is defined.

The options are:

application

Specifies that an alert is triggered for applications for which a threshold is breached.

pool-member

Specifies that an alert is triggered for pool members for which a threshold is breached.

virtual-server

Specifies that an alert is triggered for virtual servers for which a threshold is breached.

metric

Specifies a metric on which the alert is defined.

The options are:

average-page-load-time

Specifies that an alert is triggered when the average time it takes for the client to respond to a request breaches the defined threshold.

average-request-throughput

Specifies that an alert is triggered when the average number of bits per second the system processed, based on requests only, breaches the defined threshold.

average-response-throughput

Specifies that an alert is triggered when the average number of bits per second the system processed, based on responses only, breaches the defined threshold.

average-server-latency

Specifies that an alert is triggered when the average time it takes for the web server to respond to a request breaches the defined threshold.

average-tps

Specifies that an alert is triggered when the average number of transactions per second breaches the defined threshold.

max-page-load-time

Specifies that an alert is triggered when the longest time it takes for the client to respond to a request breaches the defined threshold.

max-request-throughput

Specifies that an alert is triggered when the maximum number of bits per second the system processed, based on requests only, breaches the defined threshold.

max-response-throughput

Specifies that an alert is triggered when the maximum number of bits per second the system processed, based on requests only, breaches the defined threshold.

max-server-latency

Specifies that an alert is triggered when the longest time it takes for the web server to respond to a request breaches the defined threshold.

max-tps

Specifies that an alert is triggered when the largest number of transactions per second breaches the defined threshold.

name Specifies a unique name for an alert. This option is required for the commands create, delete, and modify.

sample-period

Specifies that the alert metric is triggered when the conditions that trigger the alert last a defined amount of time, measured in seconds. The default value is 300.

threshold

Specifies the threshold that must be breached in order for the system to generate alert.

threshold-relation

Specifies whether the metric value must be below or above the metric.

The options are:

above

Specifies that an alert is issued if metric current value is above the threshold.

below

Specifies that an alert is issued if metric current value is below the threshold.

captured-traffic-external-logging

Enables or disables the external logging of captured traffic.

captured-traffic-internal-logging

Enables or disables the internal logging of captured traffic.

collect-page-load-time

Enables or disables the collection of the page load time statistics. The page load time is the round-trip latency between client end-users and the servers, that is, the round-trip time between an end-user's request for a page until the time the response finishes loading.

collect-geo

Enables or disables the collection of the names of the countries from where the traffic was sent.

collect-http-throughput

Enables or disables the collection of throughput statistics. This property has been deprecated. As of v11.3.0, HTTP throughput is always collected.

collect-http-timing-metrics

Enables or disables the collection of HTTP timing metrics.

collect-ip

Enables or disables the collection of client IPs statistics.

collect-max-tps-and-throughput

Enables or disables the collection of maximum TPS and throughput for all collected entities.

collect-methods

Enables or disables the collection of HTTP methods statistics.

`collect-response-codes`

Enables or disables the collection of response codes returned by the servers.

`collect-server-latency`

Enables or disables the collection of server latency statistics. This property has been deprecated. As of v11.3.0, server latency is always collected.

`collect-subnets`

Enables or disables the collection of client side subnets.

`collect-url`

Enables or disables the collection of requested URL statistics.

`collect-user-agent`

Enables or disables the collection of user agents.

`collect-user-sessions`

Enables or disables the collection of the unique user sessions.

`collected-stats-external-logging`

Enables or disables the external logging of the collected statistics.

`collected-stats-internal-logging`

Enables or disables the internal logging of the collected statistics.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is analytics.

`description`

User defined description.

`external-logging-publisher`

Specifies the external logging publisher used to send statistical data to one or more destinations.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`notification-by-email`

Enables or disables sending the analytics alerts by email.

`notification-by-snmp`

Enables or disables sending the analytics alerts by SNMP traps. `notification-by-syslog` must be enabled.

`notification-by-syslog`

Enables or disables logging of the analytics alerts into the Syslog.

`notification-email-addresses`

Specifies which email addresses receive alerts by email when `notification-by-email` is enabled.

`partition`

Displays the administrative partition within which the component resides.

`publish-irule-statistics`

Enables or disables publishing analytics statistics for iRules.

`sampling`

Enables or disables transaction sampling. This attribute can be set in the default profile only. The default value is disabled.

`session-cookie-security`

Specifies the condition for adding a secure attribute to the session cookie. The options are:

`always`

The secure attribute is always added to the session cookie.

`never`

The secure attribute is never added to the session cookie.

`ssl-only`

The secure attribute is only added to the session cookie when the virtual server has a client-SSL profile. This is the default value.

`session-timeout-minutes`

Specifies the number of minutes of user non-activity before the system considers the session to be over.

`smtp-config`

Specifies the SMTP configuration to be used with analytics.

`subnet-masks`

Adds, deletes, or replaces predefined subnet addresses. This options defines the display names given to certain subnet addresses seen in the client IP subnets report.

subnet

Subnet address. IPv4 addresses will be masked by 255.255.255.0. IPv6 addresses will be masked by ffff:ffff:ffff:ffff:: .

countries-for-stat-collection

Manage a list of 10 predefined countries that are used for collecting AVR statistics

ips-for-stat-collection

Manage a list of 10 predefined IP addresses that are used for collecting AVR statistics.

subnets-for-stat-collection

Manage a list of 10 predefined URLs that are used for collecting AVR statistics.

urls-for-stat-collection

Manage a list of 10 predefined subnet addresses that are used for collecting AVR statistics. IPv4 addresses will be masked by 255.255.255.0. IPv6 addresses will be masked by ffff:ffff:ffff:ffff:: .

traffic-capture

Adds, deletes, or replaces an analytics traffic capture definition. You can configure the following options for an analytics traffic capture:

app-service

Specifies the name of the application service to which the analytics traffic capture belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the analytics traffic capture. Only the application service can modify or delete the analytics traffic capture.

captured-protocols

Specifies whether the system captures traffic data sent using all protocols, or only one type of protocol.

The options are:

all Specifies that the system captures traffic data sent using all protocols.

http Specifies that the system captures traffic data sent using http protocol.

https

Specifies that the system captures traffic data sent using https protocol.

client-ips

Adds, deletes, or replaces a set of client IP addresses from/to which captured traffic is sent.

dos-activity

Specifies whether the system captures traffic data mitigated by DoS Layer 7 Enforcer or regardless of DoS activity.

The options are:

any Specifies that system does not filter traffic data by DoS activity.

mitigated-by-dosl7

Specifies that the system captures only traffic data mitigated by DoS Layer 7 Enforcer.

methods

Adds, deletes, or replaces a set of HTTP methods used to send requests from which traffic is captured.

name Specifies a unique name for an analytics traffic capture. This option is required for the commands create, delete, and modify.

node-addresses

Adds, deletes, or replaces a set of node addresses from/to which captured traffic is sent.

request-captured-parts

Specifies what parts of the request data the system captures.

The options are:

all Specifies that the system captures all the parts of the request data.

body Specifies that the system captures the body of the request data.

headers

Specifies that the system captures the HTTP headers of the request data.

none Specifies that the system does not capture the request data.

request-content-filter-search-part

Specifies which part of the request is filtered by a specific string.

The options are:

all Specifies that the system filters all the parts of the request data.

body Specifies that the system filters the body of the request data.

headers

Specifies that the system filters the HTTP headers of the request data.

none Specifies that system does not filter the request data.

uri Specifies that the system filters the URI path component, including the query string, of the request data.

request-content-filter-search-string

Specifies the string by which a request data is filtered, or none.

response-captured-parts

Specifies what parts of the response data the system captures.

The options are:

all Specifies that the system captures all the parts of the response data.

body Specifies that the system captures the body of the response data.

headers

Specifies that the system captures the HTTP headers of the response data.

none Specifies that the system does not capture the response data.

response-codes

Adds, deletes, or replaces a set of HTTP response codes from which traffic is captured.

response-content-filter-search-part

Specifies which part of the response is filtered by a specific string.

The options are:

all Specifies that the system filters all the parts of the response data.

body Specifies that the system filters the body of the response data.

headers

Specifies that the system filters the HTTP headers of the response data.

none Specifies that system does not filter the response data.

response-content-filter-search-string

Specifies the string by which the response data is filtered, or none.

url-path-prefixes

Adds, deletes, or replaces a set of URL path prefixes on which traffic can be captured (both to and from).

user-agent-substrings

Adds, deletes, or replaces a set of user agent substrings on which traffic can be captured (both to and from).

virtual-servers

Adds, deletes, or replaces a set of virtual servers from/to which captured traffic is sent.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, smtp, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2017. All rights reserved.

BIG-IP 2017-11-20 ltm profile analytics(1)

ltm profile certificate-authority

NAME

certificate-authority - Defines the settings necessary to authenticate the client certificate.

MODULE

ltm profile

SYNTAX

Configure the certificate-authority within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create certificate-authority [name]
modify certificate-authority [name]
options:
  authenticate-depth
  ca-file
  crl-file
  default-name
  description
  update-crl
```

```
edit certificate-authority [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list certificate-authority
list certificate-authority [ [name] | [glob] | [regex] ] ... ]
  app-service
  partition
```

```
show certificate-authority
show certificate-authority [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  field-fmt
  non-default-properties
  one-line
```

DESCRIPTION

Use the certificate-authority component to modify or display a certificate-authority profile.

EXAMPLES

```
create ltm profile certificate-authority mycaprofile { ca-file ca.crt }
```

Creates a certificate authority profile named mycaprofile using the system defaults.

```
modify ltm profile certificate-authority mycaprofile { authenticate-depth 3 }
```

Modifies the authenticate-depth setting to 3 for the certificate authority profile named mycaprofile.

OPTIONS

app-service

Displays the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

authenticate-depth

Specifies the authenticate depth. This is the client certificate chain maximum traversal depth.

ca-file

Specifies the certificate authority file name or, you can use default for the default certificate authority file name. Configures certificate verification by specifying a list of client or server certificate authorities that the traffic management system trusts.

crl-file

Specifies the certificate revocation list file name. You can use default for the default certificate revocation file name.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified.

description

User defined description.

name Specifies the profile instance name. This option is required for the modify command.

partition

Specifies the administrative partition within which the profile resides.

regex

Specifies the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

update-crl

Automatically updates the CRL file.

SEE ALSO

edit, glob, list, modify, regex, show, tmsh,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 ltm profile certificate-authority(1)

ltm profile classification

NAME

classification - Configures a classification profile.

MODULE

ltm profile

SYNTAX

Configure the classification profile within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create classification [name]

modify classification [name]

options:

defaults-from [[name] | none]

description [string]

app-detection [on | off]

urlcat [on | off]

irule-event [on | off]

log-publisher [string]

preset [string]

edit classification [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list classification

list classification [[name] | [glob] | [regex]] ...]

show running-config classification

show running-config classification [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete classification [name]

DESCRIPTION

Use the classification component to create, modify, delete or display a classification profile.

EXAMPLES

```
create classification my_cl { irule-event off preset /Common/ce }
```

Creates classification profile named my_cl with Classification Engine configuration /Common/ce and iRule event disabled.

```
edit classification my_cl
```

Edits the classification profile named my_cl.

```
list classification my_cl
```

Displays the properties of the my_cl classification profile.

```
delete classification my_cl
```

Deletes the my_cl classification profile.

OPTIONS

description

User defined description.

app-detection

Enables / Disables Application Detection feature.

urlcat

Enables / Disables URL categorization feature.

irule-event

Enables / disables CLASSIFICATION_DETECTED iRule event generation.

log-publisher

Specifies the log publisher name

preset

Specifies Classification Engine configuration. Refer to ltm classification ce for more details

SEE ALSO

edit, list, ltm virtual, ltm classification, ltm classification ce, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-11-12 ltm profile classification(1)

ltm profile client-ldap

NAME

client-ldap - Configures an Client LDAP profile.

MODULE

ltm profile

SYNTAX

Configure the client-ldap component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create client-ldap [name]

modify client-ldap [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

activation-mode [none | allow | require]

edit client-ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list client-ldap

list client-ldap [[[name] | [glob] | [regex]] ...]

show running-config client-ldap

show running-config client-ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DELETE

delete client-ldap [name]

DESCRIPTION

You can use the client-ldap component to create, modify, display, or delete an Client LDAP profile with which you can manage Client LDAP traffic.

EXAMPLES

```
create client-ldap my_clientldap_profile defaults-from clientldap
```

Creates a custom Client LDAP profile named my_clientldap_profile that inherits its settings from the system default Client LDAP profile.

```
list client-ldap
```

Displays the properties of all Client LDAP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot

modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is smtp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

activation-mode

Sets the activation-mode STARTTLS. The options are NONE, ALLOW, or REQUIRE. The default value is REQUIRE.

SEE ALSO

create, delete, edit, glob, list, Itm virtual, modify, regex, reset-stats, show, sys provision, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-05-06 Itm profile client-ldap(1)

Itm profile client-ssl

NAME

client-ssl - Configures a Client SSL profile.

MODULE

Itm profile

SYNTAX

Configure the client-ssl component within the Itm.profile module using the syntax shown in the following sections.

CREATE/MODIFY

create client-ssl [name]

modify client-ssl [name]

options:

alert-timeout [indefinite | immediate | [integer]]

allow-non-ssl [disabled | enabled]

allow-dynamic-record-sizing [disabled | enabled]

app-service [[string] | none]

authenticate [always | once]

authenticate-depth [integer]

bypass-on-client-cert-fail [disabled | enabled]

bypass-on-handshake-alert [disabled | enabled]

c3d-client-fallback-cert [name | none]

c3d-drop-unknown-ocsp-status [drop | ignore]

c3d-ocsp [[ocsp profile name] | none]

ca-file [name]

cache-size [integer]

cache-timeout [integer]

cert [name]

cert-extension-includes {

none |

[basic-constraints extended-key-usage

key-usage subject-alternative-name

subject-directory-attribute

]...

}

cert-key-chain [add | delete | modify | replace-all-with] {

[[name]] {

options:

cert [name | none]

chain [name | none]

key [name]

```

passphrase [none | [string] ]
usage [SERVER | CA]
}
}
cert-lifespan [integer]
cert-lookup-by-ipaddr-port [disabled | enabled]
chain [name | none]
cipher-group [name | none]
ciphers [name | none]
client-cert-ca [name | none]
crl-file [name]
allow-expired-crl [enabled | disabled]
defaults-from [clientssl | [name] ]
description [string]
destination-ip-blacklist [name]
destination-ip-whitelist [name]
forward-proxy-bypass-default-action [intercept | bypass]
generic-alert [disabled | enabled]
handshake-timeout [indefinite | [integer] ]
hostname-blacklist [name]
hostname-whitelist [name]
key [ [name] | none]
maximum-record-size [integer]
mod-ssl-methods [disabled | enabled]
mode [disabled | enabled]
notify-cert-status-to-virtual-server [disabled | enabled]
ocsp-stapling [disabled | enabled]
options {
  none |
  [ dont-insert-empty-fragments no-dtls no-dtlsv1.0 no-dtlsv1.2
no-session-resumption-on-renegotiation no-ssl no-sslv3
no-tls no-tlsv1 no-tlsv1.1 no-tlsv1.2 no-tlsv1.3 gmsslv1.1 passive-close
single-dh-use tls-rollback-bug ]...
}
passphrase [none | [string] ]
peer-cert-mode [auto | ignore | request | require]
peer-no-renegotiate-timeout [indefinite | [integer] ]
proxy-ssl [disabled | enabled]
proxy-ssl-passthrough [disabled | enabled]
proxy-ca-cert [name]
proxy-ca-key [name]
proxy-ca-lifespan [integer]
proxy-ca-passphrase [string]
renegotiate-max-record-delay [indefinite | [integer] ]
renegotiate-period [indefinite | [integer] ]
renegotiate-size [indefinite | [integer] ]
renegotiation [disabled | enabled]
retain-certificate [true | false]
secure-renegotiation [request | require | require-strict]
max-renegotiations-per-minute [integer]
max-aggregate-renegotiation-per-minute [integer]
server-name [name]
session-mirroring [disabled | enabled]
session-ticket [disabled | enabled]
session-ticket-timeout [integer]
sni-default [true | false]
sni-require [true | false]
source-ip-blacklist [name]
source-ip-whitelist [name]
ssl-c3d [disabled | enabled]
ssl-forward-proxy [disabled | enabled]
ssl-forward-proxy-bypass [disabled | enabled]
ssl-forward-proxy-verified-handshake [disabled | enabled]
strict-resume [disabled | enabled]
unclean-shutdown [disabled | enabled]
ssl-sign-hash [any | sha1 | sha256 | sha384]
max-active-handshakes [integer]
data-Ortt [disabled | enabled-with-anti-replay | enabled-no-anti-replay]

edit client-ssl [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

options:
mv client-ssl [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
to-folder

reset-stats client-ssl
reset-stats client-ssl [ [ [name] | [glob] | [regex] ] ... ]

DISPLAY
list client-ssl
list client-ssl [ [ [name] | [glob] | [regex] ] ... ]
show running-config client-ssl
show running-config client-ssl [ [ [name] | [glob] | [regex] ] ... ]
options:

```

all-properties
inherit-certkeychain
non-default-properties
one-line
partition

show client-ssl
show client-ssl [[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete client-ssl [all | [name]]
options:
recursive

DESCRIPTION

You can use the client-ssl component to create, modify, or delete a custom Client SSL profile, or display a custom or default Client SSL profile.

Client-side profiles allow the traffic management system to handle authentication and encryption tasks for any SSL connection coming into a traffic management system from a client system.

EXAMPLES

```
create client-ssl my_clientssl_profile
```

Creates a clientssl profile named my_clientssl_profile using the system defaults.

```
create clientssl my_clientssl_profile authenticate-depth number
```

Creates a Client SSL profile named my_clientssl_profile using the system defaults, except that a user is authenticated with depth number.

```
mv client-ssl /Common/my_client-ssl_profile to-folder /Common/my_folder
```

Moves a custom client-ssl profile named my_client-ssl_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

alert-timeout

Specifies the maximum time period in seconds to keep the SSL session active after alert message is sent, or indefinite. The default value is indefinite.

allow-non-ssl

Enables or disables non-SSL connections. Specify enabled when you want non-SSL connections to pass through the traffic management system as clear text. The default value is disabled.

allow-dynamic-record-sizing

Enables or disables dynamic application record sizing. Specify enabled when you want to allow dynamic record sizing. The default value is disabled.

app-service

Specifies the name of the application service to which the profile belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

authenticate

Specifies how often the system authenticates a user. The default value is once. Note that if this is set to always session cache and session ticket will be disabled.

authenticate-depth

Specifies the authenticate depth. This is the client certificate chain maximum traversal depth. The default value is 9.

bypass-on-client-cert-fail

Enables or disables SSL forward proxy bypass on failing to get client certificate that server asks for. When enabled and the SSL handshake cannot be completed because of failure to get the client certificate, SSL traffic bypasses the BIG-IP system untouched, without decryption/encryption. The default value is disabled. Conversely, you can specify enabled to use this feature.

bypass-on-handshake-alert

Enables or disables SSL forward proxy bypass on receiving handshake_failure, protocol_version or unsupported_extension alert message during the serverside SSL handshake. When enabled and there is an SSL handshake_failure, protocol_version or unsupported_extension alert during the serverside SSL handshake, SSL traffic bypasses the BIG-IP system untouched, without decryption/encryption. The default value is disabled. Conversely, you can specify enabled to use this feature.

c3d-client-fallback-cert

Specifies the client certificate to use in SSL client certificate constrained delegation. This certificate will be used if client does not provide a cert during the SSL handshake. The default value is none.

c3d-drop-unknown-ocsp-status

Specifies the BIG-IP action when the OCSP responder returns unknown status. The default value is drop, which causes the connection to be dropped. Conversely, you can specify ignore, which causes the

connection to ignore the unknown status and continue.

c3d-ocsp

Specifies the SSL client certificate constrained delegation OCSP object that the BIG-IP SSL should use to connect to the OCSP responder and check the client certificate status.

ca-file

Specifies the certificate authority (CA) file name. Configures certificate verification by specifying a list of client or server CAs that the traffic management system trusts. The default value is none.

cache-size

Specifies the SSL session cache size. For client-side profiles only, you can configure timeout and size values for the SSL session cache. Because each profile maintains a separate SSL session cache, you can configure the values on a per-profile basis. The default value is 262144.

cache-timeout

Specifies the SSL session cache timeout value. This specifies the number of usable lifetime seconds of negotiated SSL session IDs. The default value is 3600 seconds. Acceptable values are integers greater than or equal to 0 and less than or equal to 86400.

cert This option is deprecated and is maintained here for backward compatibility reasons. Please check cert-key-chain option to add certificate, key, passphrase and chain to the profile.

cert-extension-includes

Specifies the extensions of the web server certificates to be included in the generated certificates using SSL Forward Proxy. For example, { basic-constraints }. The default value is none. The extensions are:

basic-constraints

Basic Constraints are used to indicate whether the certificate belongs to a CA.

extended-key-usage

Extended Key Usage is used, typically on a leaf certificate, to indicate the purpose of the public key contained in the certificate.

key-usage

Key Usage provides a bitmap specifying the cryptographic operations which may be performed using the public key contained in the certificate; for example, it could indicate that the key should be used for signature but not for encipherment.

subject-alternative-name

Subject Alternative Name allows identities to be bound to the subject of the certificate. These identities may be included in addition to or in place of the identity in the subject field of the certificate.

subject-directory-attributes

Subject Directory Attributes are used to convey identification attributes (for example, nationality) of the subject.

destination-ip-blacklist

Specifies the data group name of destination ip blacklist when SSL forward proxy bypass feature is enabled.

destination-ip-whitelist

Specifies the data group name of destination ip whitelist when SSL forward proxy bypass feature is enabled.

forward-proxy-bypass-default-action

Specifies the SSL forward proxy bypass default action. The default option is intercept.

hostname-blacklist

Specifies the data group name of hostname blacklist when SSL forward proxy bypass feature is enabled.

hostname-whitelist

Specifies the data group name of hostname whitelist when SSL forward proxy bypass feature is enabled.

inherit-certkeychain

This is read only value used internally.

cert-key-chain

Adds, deletes, or replaces a set of certificate, key, passphrase, chain (usage specifies whether this item is used for Server or CA, where Server is the default and CA is for SSL forward proxy). client-ssl profile requires at least one cert/key pair to work. Multiple cert/key types can be associated to a client-ssl profile using following options:

cert Specifies the name of the certificate installed on the traffic management system for the purpose of terminating or initiating an SSL connection. You can specify the default certificate name, which is default.crt.

chain

Specifies or builds a certificate chain file that a client can use to authenticate the profile. The default value is none.

key Specifies the name of a key file that you generated and installed on the system. When selecting this option, type a key file name or use the default value default.key.

passphrase

Specifies the key passphrase, if required. The default value is none.

cert-lifespan

Specifies the lifespan of the certificate generated using the SSL forward proxy feature. The default value is 30.

cert-lookup-by-ipaddr-port

Specifies whether to perform certificate look up by IP address and port number.

chain

This option is deprecated and is maintained here for backward compatibility reasons. Please check cert-key-chain option to add certificate, key, passphrase and chain to the profile.

cipher-group

Specifies a cipher group. If the cipher group is not blank or none, the ciphers string will be used.

ciphers

Specifies a cipher name. The default value is DEFAULT, which uses the default ciphers.

client-cert-ca

Specifies the client cert certificate authority name. The default value is none.

crl-file

Specifies the certificate revocation list file name. The default value is none.

allow-expired-crl

Use the specified CRL file even if it has expired. The default value is disabled.

defaults-from

This setting specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is clientssl.

description

User defined description.

generic-alert

Enables or disables generic-alert. The default option is enabled, which causes the SSL profile to use generic alert number. Conversely, you can specify disabled to cause SSL profile to use alert number defined in RFC5246/RFC6066 strictly.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

handshake-timeout

Specifies the handshake timeout in seconds. The default value is 10 seconds.

key This option is deprecated and is maintained here for backward compatibility reasons. Please check cert-key-chain option to add certificate, key, passphrase and chain to the profile.

maximum-record-size

Specifies the profile's maximum record size. The range is 128 - 16384. The default value is 16384.

mod-ssl-methods

Enables or disables ModSSL method emulation. Enable this option when OpenSSL methods are inadequate, for example, when you want to use SSL compression over TLSv1. The default value is disabled.

mode Specifies the profile mode, which enables or disables SSL processing. The default value is enabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

options

Enables options, including some industry-related workarounds. Enter options inside braces, for example, {dont-insert-empty-fragments}.

The default value is dont-insert-empty-fragments no-tlsv1.3. The options are:

dont-insert-empty-fragments

Disables a countermeasure against an SSL 3.0/TLS 1.0 protocol vulnerability affecting CBC ciphers. These ciphers cannot be handled by certain broken SSL implementations. This option has no effect for connections using other ciphers.

no-session-resumption-on-renegotiation

When performing renegotiation as an SSL server, this option always starts a new session (that is, session resumption requests are only accepted in the initial handshake). The system ignores this option for server-side SSL.

gmsslv1.1

Enable GMSSLv1.1 protocol.

no-ssl

Do not use any version of the SSL protocol.

no-sslv3

Do not use the SSLv3 protocol.

no-tls

Do not use any version of the TLS protocol.

no-tlsv1

Do not use the TLSv1.0 protocol.

no-tlsv1.1

Do not use the TLSv1.1 protocol.

no-tlsv1.2

Do not use the TLSv1.2 protocol.

no-tlsv1.3

Do not use the TLSv1.3 protocol.

no-dtls

Do not use any version of the DTLS protocol.

no-dtlsv1.0

Do not use the DTLSv1.0 protocol.

no-dtlsv1.2

Do not use the DTLSv1.2 protocol.

passive-close

Specifies how to handle passive closes.

none Disables all workarounds. Note that F5 Networks does not recommend this option.

notify-cert-status-to-virtual-server

Specifies whether to propagate the status of the certificates of this clientssl profile to the virtual servers that are using this clientssl profile.

ocsp-stapling

Specifies whether to enable OCSP stapling.

single-dh-use

Creates a new key when using temporary/ephemeral DH parameters. This option must be used to prevent small subgroup attacks, when the DH parameters were not generated using strong primes (for example, when using DSA-parameters). If strong primes were used, it is not strictly necessary to generate a new DH key during each handshake, but F5 Networks recommends it. Enable the Single DH Use option whenever temporary or ephemeral DH parameters are used.

tls-rollback-bug

Disables version rollback attack detection. During the client key exchange, the client must send the same information about acceptable SSL/TLS protocol levels as it sends during the first hello. Some clients violate this rule by adapting to the server's answer. For example, the client sends an SSLv2 hello and accepts up to SSLv3.1 (TLSv1), but the server only processes up to SSLv3. In this case, the client must still use the same SSLv3.1 (TLSv1) announcement. Some clients step down to SSLv3 with respect to the server's answer and violate the version rollback protection. The system ignores this option for server-side SSL.

partition

Displays the administrative partition within which the profile resides.

passphrase

This option is deprecated and is maintained here for backward compatibility reasons. Please check cert-key-chain option to add certificate, key, passphrase and chain to the profile.

peer-cert-mode

Specifies the peer certificate mode. The default value is ignore.

peer-no-renegotiate-timeout Specifies the timeout in seconds when the server sends Hello Request and waits for ClientHello before it sends Alert with fatal alert. You can also specify indefinite. The default is 10 seconds.

proxy-ca-cert

Specifies the name of the certificate file that is used as the certification authority certificate when SSL forward proxy feature is enabled. The certificate should be generated and installed by you on the system. When selecting this option, type a certificate file name. (This option is deprecated since v14.0.0, suggest to use cert-key-chain with usage CA to add SSL forward proxy CA key/cert.)

proxy-ca-key

Specifies the name of the key file that is used as the certification authority key when SSL forward proxy feature is enabled. The key should be generated and installed by you on the system. When selecting this option, type a key file name. (This option is deprecated since v14.0.0, suggest to use cert-key-chain with usage CA to add SSL forward proxy CA key/cert.)

proxy-ca-passphrase

Specifies the passphrase of the key file that is used as the certification authority key when SSL forward proxy feature is enabled. When selecting this option, type the passphrase corresponding to the selected proxy-ca-key. (This option is deprecated since v14.0.0, suggest to use cert-key-chain with usage CA to add SSL forward proxy CA key/cert.)

proxy-ssl

Enabling this option requires a corresponding server ssl profile with proxy-ssl enabled to perform transparent SSL decryption. This allows further modification of application traffic within an SSL tunnel while still allowing the server to perform necessary authorization, authentication, auditing steps.

proxy-ssl-passthrough

Enabling this option requires a corresponding server ssl profile with proxy-ssl-passthrough enabled. This allows Proxy SSL to passthrough the traffic when ciphersuite negotiated between the client and server is not supported. The default option is disabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

renegotiate-max-record-delay

Specifies the maximum number of SSL records that the traffic management system can receive before it renegotiates an SSL session. After the system receives this number of SSL records, it closes the connection. This setting applies to client profiles only. The default value is indefinite.

renegotiate-period

Specifies the number of seconds required to renegotiate an SSL session. The default value is indefinite.

renegotiate-size

Specifies the size of the application data, in megabytes, that is transmitted over the secure channel. If the size of the data is higher than this value, the traffic management system must renegotiate the SSL session. The default value is indefinite.

renegotiation

Specifies whether renegotiations are enabled. The default value is enabled. When renegotiations are disabled, and the system is acting as an SSL server, and a COMPAT or NATIVE cipher is negotiated, the system will abort the connection. Additionally, when renegotiations are disabled, and the system is acting as an SSL client, the system will ignore the server's HelloRequest messages.

retain-certificate

APM module requires storing certificate in SSL session. When set to false, certificate will not be stored in SSL session. The default value is true.

secure-renegotiation

Specifies the secure renegotiation mode. The default value is require. When secure renegotiation is required, any client attempting to renegotiate that does not support secure renegotiation will have its connection aborted. When secure renegotiation is set to require-strict, any client attempting to connect that does not support secure renegotiation will have its initial handshake denied. When secure renegotiation is set to request, unpatched clients will be permitted to renegotiate. This setting is NOT recommended however, as it is subject to active man-in-the-middle attacks.

max-renegotiations-per-minute

Specifies the maximum number of renegotiation attempts allowed in a minute. The default value is 5.

max-active-handshakes

Specifies the maximum number allowed SSL active handshakes. The default value is 0.

max-aggregate-renegotiation-per-minute

Specifies the maximum number of aggregate renegotiation attempts allowed in a minute. The default value is indefinite.

server-name

Specifies the server names to be matched with SNI (server name indication) extension information in ClientHello from a client connection. Wildcard is supported by using wildcard character "*" to match multiple names.

sni-default

When true, this profile is the default SSL profile when the server name in a client connection does not match any configured server names, or a client connection does not specify any server name at all.

sni-require

When this option is enabled, a client connection that does not specify a known server name or does not support SNI extension will be rejected.

ssl-sign-hash

Specifies SSL sign hash algorithm which is used to sign and verify SSL Server Key Exchange and Certificate Verify messages for the specified SSL profiles. The default value is sha1.

strict-resume

Enables or disables strict-resume. The default option is disabled, which causes the SSL profile to resume an uncleanly shut down SSL session. Conversely, you can specify enabled to prevent an SSL session from being resumed after an unclean shutdown.

unclean-shutdown

By default, the SSL profile performs unclean shutdowns of all SSL connections, which means that underlying TCP connections are closed without exchanging the required SSL shutdown alerts. If you want to force the SSL profile to perform a clean shutdown of all SSL connections, set this option to disabled.

session-mirroring

Enables or disables the mirroring of sessions to high availability peer. By default, this setting is disabled, which causes the system to not mirror ssl sessions.

session-ticket

Enables or disables session-ticket. The default option is disabled, which causes the SSL profile not to use session ticket per RFC 5077. Conversely, you can specify enabled to cause SSL profile to use session ticket per RFC 5077.

session-ticket-timeout

Specifies the session ticket timeout. The default value is 0 which means cache timeout is used.

source-ip-blacklist

Specifies the data group name of source ip blacklist when SSL forward proxy bypass feature is enabled.

source-ip-whitelist

Specifies the data group name of source ip whitelist when SSL forward proxy bypass feature is enabled.

data-0rtt

Specifies if TLSv1.3 should accept 0-RTT with early data, with or without anti-replay. To protect against packet replay, F5 recommends that you enable anti-replay. The default value is disabled, which means TLSv1.3 will discard any early data.

ssl-c3d

Enables or disables SSL client certificate constrained delegation. The default option is disabled. Conversely, you can specify enabled to use the SSL client certificate constrained delegation.

ssl-forward-proxy

Enables or disables SSL forward proxy feature. The default option is disabled. Conversely, you can specify enabled to use the SSL Forward Proxy Feature.

ssl-forward-proxy-bypass

Enables or disables SSL forward proxy bypass feature. The default option is disabled. Conversely, you can specify enabled to use the SSL Forward Proxy Bypass Feature.

ssl-forward-proxy-verified-handshake

Specifies, when enabled, that in SSL forward proxy mode, the system should always do a TLS handshake with the server first before doing the client handshake. When disabled, the system will do the server handshake first only if it has not previously forged and cached the server certificate; once the server certificate is ready, the system will always handshake first with the client. The default value is disabled.

to-folder

client-ssl profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2019-12-20 ltm profile client-ssl(1)

ltm profile dhcpv4

NAME

DHCPv4 - Configures a Dynamic Host Configuration Protocol (DHCP) profile.

MODULE

ltm profile

SYNTAX

Configure the dhcpv4 profile within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create dhcpv4 [name]

modify dhcpv4 [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

mode [relay | forwarding]

description [string]

idle-timeout [[integer] | indefinite]

default-lease-time [integer]

lease-query-max-retry [integer]

lease-query-only [true | false]

transaction-timeout [integer]

authentication {

options:

enabled [true | false]

virtual [[string] | none]

user-name {

options:

```

format [mac-address | mac-and-relay-option | relay-option | tcl-snippet]
suboption-id1 [integer]
suboption-id2 [integer]
separator1 [[string] | none]
separator2 [[string] | none]
tcl [[string] | none]
}
}
subscriber-discovery {
  options:
enabled [true | false]
  subscriber-id {
    options:
format [mac-address | mac-and-relay-id | tcl-snippet]
suboption-id1 [integer]
suboption-id2 [integer]
separator1 [[string] | none]
separator2 [[string] | none]
tcl [[string] | none]
}
}
}
relay-agent-id {
  options:
add [true | false]
remove [true | false]
suboption {
  options:
id1 [integer]
id2 [integer]
value1 [string | none]
value2 [string | none]
}
}
}
ttl-value [integer]
ttl-dec-value [ by-0 | by-1 | by-2 | by-4 ]
max-hops [integer]

```

```

edit dhcpv4 [ [name] ... ]
options:
  all-properties
  non-default-properties

```

```

reset-stats dhcpv4
reset-stats dhcpv4 [ [ [name] | [regex] ] ... ]

```

```

DISPLAY
list dhcpv4
list dhcpv4 [ [ [name] | [regex] ] ... ]
show running-config dhcpv4
show running-config dhcpv4
[ [ [name] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

```

```

show dhcpv4
show dhcpv4 [ [ [name] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global

```

```

DELETE
delete dhcpv4 [name]

```

DESCRIPTION
You can use the dhcpv4 profile to manage DHCPv4 network traffic.

EXAMPLES
create dhcpv4 my_dhcpv4_profile defaults-from dhcpv4

Creates a custom DHCPv4 profile named my_dhcpv4_profile that inherits its settings from the system default DHCPv4 profile.

list dhcpv4 all-properties

Displays all properties for all DHCPv4 profiles.

```

create dhcpv4 new_dhcpv4_profile {
mode relay
idle-timeout 120
transaction-timeout 45
default-lease-time 3600
ttl-value 0
ttl-dec-value by-2

```

```

max-hops 4
subscriber-discovery {
    enabled true
    subscriber-id {
    format mac-and-relay-option
    suboption-id1 1
    separator1 -
    }
}
authentication {
    enabled true
    virtual new_authen_vs
    user-name {
    format mac-and-relay-option
    suboption-id1 1
    separator1 -
    }
}
relay-agent-id {
    add false
    remove true
} }

```

Creates a DHCPv4 profile named `new_dhcpv4_profile` with `idle-timeout` value of `<120 seconds>`, `transaction-timeout` of `<45 seconds>`, `default-lease-time` of `3600 seconds`, `tll-dec-value` of `2` and `max-hops` of `4`. The BIG-IP virtual will work in relay mode, with subscriber discovery enabled and configured with the `subscriber-id` format set to `mac-and-relay-option` and configured to use only suboption 1 (the first suboption ID) of the relay-agent info option (option 82) and to use the `-` to concatenate the MAC address to the first suboption ID. The authentication is enabled and its `user-name` equals the `subscriber-id` and the authentication virtual name is `new_authen_vs`. It also does not add relay agent option 82 but removed it (if exists) from the server-to-client messages.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.
 Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

default-lease-time

Provides the default value in seconds of DHCPv4 lease time in case it was missing in the client-server exchange. The default is 86400.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is `dhcpv4`.

description

User defined description.

lease-query-max-retry

Specifies the maximum number of retries for a DHCPv4 lease query request.

lease-query-only

Specifies if a DHCPv4 virtual that uses this profile can be used for DHCP lease query only or for DHCP discovery as well.

`mode` Specifies the operation mode of the DHCP virtual. If the virtual to run in relay mode, then it means that it is acting as a standard DHCPv4 relay agent. This means that the relay will change some of the DHCPv4 packet fields before sending it to either the client or server. The forwarding mode is similar to relay except that the virtual will not modify the standard fields, instead it will forward the message from client to server and vice-versa. The default is `relay`.

idle-timeout

Specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 60 seconds.

transaction-timeout

Specifies DHCPv4 transaction timeout, in seconds. The transactions should complete within the timeout specified. If a transaction does not complete for any reason, it is removed. The default value is 45 seconds.

authentication

Manages the subscriber authentication attributes.

enabled

To enable or disable subscriber authentication. If enabled, then user has to fill the following fields. Default is false.

virtual

Specifies the authentication virtual server name.

user-name

Manages the authentication user name's attributes. The `user-name` is what will be used to authenticate the DHCP client.

format

Specifies the `user-name` format. The options are: MAC address, MAC + relay-agent option, relay-

agent option or tcl-snippet. The concatenation symbol is defined as separator1.

suboption-id1

The relay-agent option (option 82) first suboption ID. The default is 1.

suboption-id2

The relay-agent option (option 82) second suboption ID. The default is 2.

separator1

A string that is used to concatenate the MAC address and the relay-agent info option (option 82) to create the authentication user-name. The default is @.

separator2

A string that is used to concatenate the relay-agent info option (option 82) suboptions 1 and 2 to create the authentication user-name. The default is @.

tcl-snippet

A tcl snippet to format the user name. This value will be taken into account only if the format value was chosen to be tcl-snippet.

subscriber-discovery

Manages the subscriber discovery attributes.

enabled

To enable or disable subscriber discovery. If enabled, then user has to fill the following fields. Default is false.

subscriber-id

Manages the subscriber-id attributes. The subscriber-id is used by SPM to create, delete and update subscriber sessions.

format

Specifies the subscriber-id format. The options are: MAC address, MAC + relay-agent option, relay-agent option or tcl-snippet. The concatenation symbol is defined as separator1.

suboption-id1

The relay-agent info option (option 82) first suboption ID. The default is 1.

suboption-id2

The relay-agent info option (option 82) second suboption ID. The default is 2.

separator1

A string that is used to concatenate the MAC address and the relay-agent info option (option 82) to create the subscriber-id. The default is @.

separator2

A string that is used to concatenate the relay-agent info option (option 82) suboptions 1 and 2 to create the subscriber-id. The default is @.

tcl-snippet

A tcl snippet to format the subscriber-id. This value will be taken into account only if the format value was chosen to be tcl-snippet.

relay-agent-id

Manages the relay agent information option (option 82) attributes. As a relay, the DHCP virtual can insert this option.

add Specifies if the user wants the DHCP relay agent to insert option 82 or not. Default is false.

remove

Specifies if the user wants the DHCP relay agent to remove option 82 from the server-to-client traffic or not. Default is false.

suboptions

Manages the inserted relay agent information option (option 82) suboptions. We allow only two suboptions to be inserted.

id1 An integer to represent the first suboption ID. Default is 1.

value1

A string to represent the first suboption value.

id2 An integer to represent the second suboption ID. Default is 2.

value2

A string to represent the second suboption value.

ttl-value

Specifies the ttl absolute value that the user may want to set for each outgoing DHCP packet. Default is 0; and in this case, we use the ttl-dec-value field.

ttl-dec-value

Specifies the amount that the DHCP virtual will use to decrement the ttl for each outgoing DHCP packet. Default is by-1.

max-hops

Specifies the maximum number of relay agents that the DHCPv4 messages pass through before they are

discarded. The default is 4.

SEE ALSO

create, delete, edit, ltm profile, ltm virtual, modify, show, dhcpv6 profile, reset-stats, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014. All rights reserved.

BIG-IP 2017-07-21 ltm profile dhcpv4(1)

ltm profile dhcpv6

NAME

DHCPv6 - Configures a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) profile.

MODULE

ltm profile

SYNTAX

Configure the dhcpv6 profile within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create dhcpv6 [name]

modify dhcpv6 [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

mode [relay | forwarding]

description [string]

idle-timeout [[integer] | indefinite]

default-lease-time [integer]

lease-query-max-retry [integer]

lease-query-only [true | false]

transaction-timeout [integer]

authentication {

options:

enabled [true | false]

virtual [[string] | none]

user-name {

options:

format [mac-address | option37 | mac-and-option37 | option38 | mac-and-option38 | option37-and-option38 | mac-and-option37-and-option38]

separator1 [[string] | none]

separator2 [[string] | none]

tcl [[string] | none]

}

}

subscriber-discovery {

options:

enabled [true | false]

subscriber-id {

options:

format [mac-address | option37 | mac-and-option37 | option38 | mac-and-option38 | option37-and-option38 | mac-and-option37-and-option38]

separator1 [[string] | none]

separator2 [[string] | none]

tcl [[string] | none]

}

}

remote-id-option {

add [true | false]

remove [true | false]

enterprise-number [integer]

value [string | none]

}

subscriber-id-option {

add [true | false]

remove [true | false]

value [string | none]

}

edit dhcpv6 [[name] ...]

options:

all-properties

non-default-properties

reset-stats dhcpv6

```
reset-stats dhcpv6 [ [ [name] | [regex] ] ... ]
```

DISPLAY

```
list dhcpv6
```

```
list dhcpv6 [ [ [name] | [regex] ] ... ]
```

```
show running-config dhcpv6
```

```
show running-config dhcpv6
```

```
[ [ [name] | [regex] ] ... ]
```

options:

- all-properties

- non-default-properties

- one-line

- partition

```
show dhcpv6
```

```
show dhcpv6 [ [ [name] | [regex] ] ... ]
```

options:

- (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

- field-fmt

- global

DELETE

```
delete dhcpv6 [name]
```

DESCRIPTION

You can use the dhcpv6 profile to manage DHCPv6 network traffic.

EXAMPLES

```
create dhcpv6 my_dhcpv6_profile defaults-from dhcpv6
```

Creates a custom DHCPv6 profile named my_dhcpv6_profile that inherits its settings from the system default DHCPv6 profile.

```
list dhcpv6 all-properties
```

Displays all properties for all DHCPv6 profiles.

```
create dhcpv6 new_dhcpv6_profile {
```

```
mode relay
```

```
idle-timeout 120
```

```
transaction-timeout 45
```

```
default-lease-time 3600
```

```
subscriber-discovery {
```

```
enabled true
```

```
subscriber-id {
```

```
format mac-and-option37
```

```
separator1 @
```

```
}
```

```
}
```

```
authentication {
```

```
enabled true
```

```
virtual new_authen_vs
```

```
user-name {
```

```
format mac-and-option37
```

```
separator1 @
```

```
}
```

```
}
```

```
remote-id-option {
```

```
add false
```

```
remove true
```

```
}
```

```
subscriber-id-option {
```

```
add false
```

```
remove true
```

```
}
```

```
}
```

Creates a DHCPv6 profile named new_dhcpv6_profile with idle-timeout value of <120 seconds>, transaction-timeout of <45 seconds> and default-lease-time of 3600 seconds. The BIG-IP virtual will work in relay mode, with subscriber discovery enabled and configured with the subscriber-id format set to mac-and-option37 (remote-id relay agent option) and to use the @ to concatenate both MAC address and option 37. The authentication is enabled and its user-name equals the subscriber-id and the authentication virtual name is new_authen_vs. It also does not add either option 37 or options 38 (remote-id option and subscriber-id option) but remove them (if exists) from the server-to-client messages.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

default-lease-time

Provides the default value in seconds of DHCPv6 lease time in case it was missing in the client-server exchange. The default is 86400.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is dhcpv6.

description

User defined description.

lease-query-max-retry

Specifies the maximum number of retries for a DHCPv6 lease query request.

lease-query-only

Specifies if a DHCPv6 virtual that uses this profile can be used for DHCP lease query only or for DHCP discovery as well.

mode Specifies the operation mode of the DHCP virtual. If the virtual to run in relay mode, then it means that it is acting as a standard DHCPv6 relay agent. This means that the relay will encapsulate the original messages into one of the relay messages before it send it to the server or the client. In the forwarding mode, the virtual will just forward the message to either the server or the client. The default is relay.

idle-timeout

Specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 60 seconds.

transaction-timeout

Specifies DHCPv6 transaction timeout, in seconds. The transactions should complete within the timeout specified. If a transaction does not complete for any reason, it is removed. The default value is 45 seconds.

authentication

Manages the subscriber authentication attributes.

enabled

To enable or disable subscriber authentication. If enabled, then user has to fill the following fields. Default is false.

virtual

Specifies the authentication virtual server name.

user-name

Manages the authentication user name's attributes. The user-name is what will be used to authenticate the DHCP client.

format

Specifies the user-name format. The options are: MAC address, option37, MAC + option37, option38, MAC + option38, option37 + option38, mac + option37 + option38 tcl-snippet. The concatenation symbols are defined as separator1 and separator2.

separator1

A string that is used to concatenate the MAC address and the first two strings of the authentication user-name. The default is @.

separator2

A string that is used to concatenate the MAC address and the second two strings of the authentication user-name. The default is @.

tcl-snippet

A tcl snippet to format the user name. This value will be taken into account only if the format value was chosen to be tcl-snippet.

subscriber-discovery

Manages the subscriber discovery attributes.

enabled

To enable or disable subscriber discovery. If enabled, then user has to fill the following fields. Default is false.

subscriber-id

Manages the subscriber-id attributes. The subscriber-id is used by SPM to create, delete and update subscriber sessions.

format

Specifies the user-name format. The options are: MAC address, option37, MAC + option37, option38, MAC + option38, option37 + option38, mac + option37 + option38 tcl-snippet. The concatenation symbols are defined as separator1 and separator2.

separator1

A string that is used to concatenate the first two strings of the subscriber-id. The default is @.

separator2

A string that is used to concatenate the second two strings of the subscriber-id. The default is @.

tcl-snippet

A tcl snippet to format the subscriber-id. This value will be taken into account only if the format value was chosen to be tcl-snippet.

remote-id-option

Manages the DHCPv6 relay agent remote-id option (option 37) attributes. As a relay, the DHCP virtual can insert or remove this option.

add Specifies if the user wants the DHCP relay agent to insert option 37 or not. Default is false.

remove

Specifies if the user wants the DHCP relay agent to remove option 37 from the server-to-client traffic or not. Default is false.

enterprise-number

Specifies the enterprise number of the inserted remote-id option (option 37).

value

A string to represent the remote-id option value.

subscriber-id-option

Manages the DHCPv6 relay agent subscriber-id option (option 38) attributes. As a relay, the DHCP virtual can insert or remove this option.

add Specifies if the user wants the DHCP relay agent to insert option 38 or not. Default is false.

remove

Specifies if the user wants the DHCP relay agent to remove option 38 from the server-to-client traffic or not. Default is false.

value

A string to represent the subscriber-id option value.

SEE ALSO

create, delete, edit, ltm profile, ltm virtual, modify, show, dhcpv6 profile, reset-stats, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2015. All rights reserved.

BIG-IP 2017-07-21 ltm profile dhcpv6(1)

ltm profile diameter

NAME

diameter - Configures a profile to manage Diameter network traffic.

MODULE

ltm profile

SYNTAX

Configure the diameter component within the ltm profile module using the syntax in the following sections.

CREATE/MODIFY

create diameter [name]

modify diameter [name]

options:

app-service [[string] | none]

connection-prime [disabled | enabled]

defaults-from [name]

description [string]

destination-realm [string]

handshake-timeout [number]

host-ip-rewrite [disabled | enabled]

max-retransmit-attempts [number]

max-watchdog-failure [number]

origin-host-to-client [string]

origin-host-to-server [string]

origin-realm-to-client [string]

origin-realm-to-server [string]

overwrite-destination-host [disabled | enabled]

parent-avp [[number] | [string]]

persist-avp [[number] | [string]]

reset-on-timeout [disabled | enabled]

retransmit-timeout [number]

watchdog-timeout [number]

mv diameter [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

edit diameter [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats diameter

reset-stats diameter [[[name] | [glob] | [regex]] ...]

DISPLAY

list diameter

list diameter [[[name] | [glob] | [regex]] ...]

show running-config diameter

show running-config diameter [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show diameter

show diameter [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete diameter [name]

DESCRIPTION

You can use the diameter component to configure a profile to manage Diameter network traffic.

EXAMPLES

```
create diameter my_diameter_profile defaults-from diameter
```

Creates a Diameter profile named `my_diameter_profile` that inherits its settings from the system default Diameter profile.

```
list diameter
```

Displays the properties of all Diameter profiles.

```
mv diameter /Common/my_diameter_profile to-folder /Common/my_folder
```

Moves a custom diameter profile named `my_diameter_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

connection-prime

When enabled, and the system receives a capabilities exchange request from the client, the system will establish connections and perform handshaking with all the servers prior to sending the capabilities exchange answer to the client. The default value is disabled.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is diameter.

description

User defined description.

destination-realm

This attribute has been deprecated as of BIG-IP v11.3.0. Specifies the realm to which messages are routed. A value of none indicates that the destination-realm option is disabled. The default value is none.

You can specify a fully qualified domain name as an ASCII string. For more information about this option, see RFC 3588 section 6.6.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

handshake-timeout

Specifies the handshake timeout in seconds. This setting specifies the maximum number of seconds that a connection can be idle after the capabilities exchange request was sent to the server. The default value is 10. The system will reset the connection after it has timed out.

You can specify a numeric value in the range 0 to 4294967295. The recommended value is in the range of 5 to 30

host-ip-rewrite

When enabled and the message is a capabilities exchange request or capabilities exchange answer, rewrite the host-ip-address attribute with the system's egress IP address. The default value is enabled.

max-retransmit-attempts

Specifies the maximum number of retransmit attempts. This setting specifies the maximum number of

attempts that BIG-IP will take to retransmit the request messages if it does not receive the corresponding answer messages. If retransmit is unsuccessful, after maximum attempts, BIG_IP will send an error response. The default value is 1.

You can specify a numeric value in the range 0 to 4294967295. The recommended value is in the range of 1 to 10

max-watchdog-failure

Specifies the maximum number of device watchdog failures that the traffic management system can take before it tears down the connection. After the system receives this number of device watchdog failures, it closes the connection. The default value is 10.

You can specify a numeric value in the range 0 to 4294967295.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

origin-host

This attribute has been deprecated as of BIG-IP v11.3.0. Please use, origin-host-to-client or origin-host-to-server. Specifies the origin host of BIG-IP. The origin-host is used to overwrite the server's actual origin host attribute when it responds to the client. A value of none indicates that origin-host is disabled. The default value is none.

You can specify an ASCII string as a FQDN. See RFC 3588 section 6.3.

origin-host-to-client

Specifies the origin host to client of BIG-IP. The origin-host-to-client is used to overwrite the server's actual origin host attribute when it responds to the client. A value of none indicates that origin-host-to-client is disabled. The default value is none.

You can specify an ASCII string as a FQDN. See RFC 3588 section 6.3.

origin-host-to-server

Specifies the origin host to server of BIG-IP. The origin-host-to-server is used to overwrite the client's actual origin host attribute when it responds to the server. A value of none indicates that origin-host-to-server is disabled. The default value is none.

You can specify an ASCII string as a FQDN. See RFC 3588 section 6.3.

origin-realm

This attribute has been deprecated as of BIG-IP v11.3.0. Please use, origin-realm-to-client or origin-realm-to-server. Specifies the origin realm of BIG-IP. The origin-realm is used to overwrite the server's actual origin realm attribute when it responds to the client. A value of none indicates that origin-realm is disabled. The default value is none.

You can specify an ASCII string as a FQDN. See RFC 3588 section 6.4.

origin-realm-to-client

Specifies the origin realm of BIG-IP. The origin-realm-to-client is used to overwrite the server's actual origin realm attribute when it responds to the client. A value of none indicates that origin-realm-to-client is disabled. The default value is none.

You can specify an ASCII string as a FQDN. See RFC 3588 section 6.4.

origin-realm-to-server

Specifies the origin realm to server of BIG-IP. The origin-realm-to-server is used to overwrite the client's actual origin realm attribute when it responds to the server. A value of none indicates that origin-realm-to-server is disabled. The default value is none.

You can specify an ASCII string as a FQDN. See RFC 3588 section 6.4.

overwrite-destination-host

This attribute has been deprecated as of BIG-IP v11.3.0. When you enable this option, the system replaces the value of the destination host field in the Diameter header with the BIG-IP(r) pool member address. When you disable this option, the system does not modify the destination host field. The default value is enabled.

parent-avp

Specifies the name of the Diameter attribute that the system uses to indicate if the persist-avp option is embedded in a grouped avp. A value of none indicates that the value of the persist-avp option is not embedded in a grouped avp. The default value is none.

You can specify an ASCII string or a numeric ID in the range 1 to 4294967295. Acceptable strings can be found in RFC 3588 section 4.5.

partition

Displays the administrative partition within which the profile resides.

persist-avp

Specifies the name of the Diameter attribute that the system persists on. A value of none indicates that persistence is disabled. The default value is session-id.

You can specify an ASCII string or a numeric ID in the range 1 to 4294967295. Acceptable strings can be found in RFC 3588 section 4.5.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at

sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reset-on-timeout

When it is enabled and the watchdog failures exceed the max watchdog failure, the system resets the connection. The default value is enabled.

retransmit-timeout

Specifies the retransmit timeout in seconds. This setting specifies the number of seconds to retransmit the request messages if BIG-IP does not receive the corresponding answer messages . The default value is 10.

You can specify a numeric value in the range 0 to 4294967295. The recommended value is in the range of 5 to 30

to-folder

diameter profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

watchdog-timeout

Specifies the watchdog timeout in seconds. This setting specifies the number of seconds that a connection is idle before the device watchdog request is sent. The default value is 0, which means BIG-IP will not send a device watchdog request to either client or server side.

You can specify a numeric value in the range 0 to 4294967295. The recommended value is in the range of 6 to 30

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 ltm profile diameter(1)

ltm profile dns-logging

NAME

dns-logging - Configures a domain name service logging (DNS Logging) profile.

MODULE

ltm profile

SYNTAX

Configure the dns-logging component within the ltm profile module using the syntax in the following sections.

CREATE/MODIFY

create dns-logging [name]

modify dns-logging [name]

options:

description [string]

enable-query-logging [no | yes]

enable-response-logging [no | yes]

include-complete-answer [no | yes]

include-query-id [no | yes]

include-source [no | yes]

include-timestamp [no | yes]

include-view [no | yes]

log-publisher [name]

edit dns-logging [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv dns-logging [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list dns-logging

list dns-logging [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line
partition

DELETE
delete dns-logging [name]

DESCRIPTION

You can use this component to create, modify, display, or delete a DNS logging profile, to enable query or response logging, and to define the format of messages themselves.

EXAMPLES

list dns-logging

Displays the properties of all DNS logging profiles.

create dns-logging my_dns_log_profile enable-query-logging yes log-publisher my_pub include-query-id yes

Creates a DNS logging profile with query logging enabled. Messages will be sent to publisher my_pub. Messages will contain the query ID.

mv dns-logging /Common/my_dns_logging_profile to-folder /Common/my_folder

Moves a custom dns-logging profile named my_dns_logging_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

description
User defined description.

enable-query-logging
Log the contents of DNS queries. The default value for this option is yes.

enable-response-logging
Log the contents of DNS responses. The default value is no.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

include-complete-answer
Selects whether all the resource records are included in response log messages. The default value is yes (complete-answer).

include-query-id
Selects whether the query id sent by the client is included in the query and response log messages. The default value is no.

include-source
Selects whether the message originator is included in the query and response log messages. The default value is yes.

include-timestamp
Selects whether the time stamp of the message is included in the query and response log messages. The default value is yes. You may or may not need this depending on whether the destination log servers prepend a time stamp to messages.

include-view
Selects whether the view is included in the query log messages. The default value is yes.

log-publisher
Specifies the log publisher used to deliver messages to one or more destinations. This option must be specified.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

to-folder
dns-logging profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, modify, mv, regex, ltm profile dns, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2015-07-14 ltm profile dns-logging(1)

Itm profile dns

NAME

dns - Configures a Domain Name System (DNS) profile.

MODULE

itm profile

SYNTAX

Configure the dns component within the Itm profile module using the syntax in the following sections.

CREATE/MODIFY

create dns [name]

modify dns [name]

options:

app-service [(string) | none]

avr-dnsstat_sample_rate [integer]

cache [string]

defaults-from [[name] | none]

description [string]

dns64 [disabled | secondary | immediate | v4-only]

dns64-additional-section-rewrite [disabled | v6-only | v4-only | any]

dns64-prefix [IPv6 prefix]

dns-security [string]

edns0-client-subnet-insert [disabled | enabled]

enable-cache [no | yes]

enable-dnssec [no | yes]

enable-dns-express [no | yes]

enable-dns-firewall [no | yes]

enable-gtm [no | yes]

enable-hardware-query-validation [no | yes]

enable-hardware-response-cache [no | yes]

enable-logging [no | yes]

enable-rapid-response [no | yes]

log-profile [[name] | none]

process-rd [no | yes]

process-xfr [no | yes]

rapid-response-last-action [allow | drop | noerror | nxdomain | refuse | truncate]

unhandled-query-action [allow | drop | hint | noerror | reject]

use-local-bind [no | yes]

edit dns [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv dns [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats dns

reset-stats dns [[[name] | [glob] | [regex]] ...]

DISPLAY

list dns

list dns [[[name] | [glob] | [regex]] ...]

show running-config dns

show running-config dns [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete dns [name]

DESCRIPTION

You can use this component to create, modify, display, or delete a DNS profile to define how the BIG-IP system handles DNS traffic. You can also display and reset DNS profile statistics.

EXAMPLES

```
create dns my_dns_profile defaults-from dns
```

Creates a DNS profile named my_dns_profile that inherits its settings from the system default DNS profile.

```
list dns
```

Displays the properties of all DNS profiles.

```
mv dns /Common/my_dns_profile to-folder /Common/my_folder
```

Moves a custom dns profile named my_dns_profile to a folder named my_folder, where my_folder has already been

created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

avr-dnsstat-sample-rate

Sets AVR DNS statistics rate. The default value is 0, which means AVR DNS statistics is disabled. If the sampling rate is set to 1, each query will be sent to the analytics database. If the sampling rate is set to an integer N, every Nth query will be sent and the analytics database will count it N times. When sampling rate is greater than one, the statistics will be inaccurate if the traffic volume is low. However, when the traffic volume is high, the system performance will benefit from sampling and the inaccuracy will be negligible. Also be aware that analytics database has its own internal sampling mechanism. The sampling rate does not apply to DNS firewall statistics. AVR DNS statistics contain query name, query type, virtual server IP and client IP.

cache

Specifies the user-created cache that the system uses to cache DNS responses. When you select a cache for the system to use, you must also enable the DNS cache setting.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is dns.

description

User defined description.

dns64

Sets DNS64 mapping mode. The default value is disabled.

dns64-additional-section-rewrite

Sets DNS64 additional section rewriting. For AAAA and A records in additional section, this field specifies how they are being rewritten. The default value is disabled.

dns64-prefix

Specifies DNS64 mapping IPv6 prefix.

dns-security

Indicates the DNS security profile the system uses.

edns0-client-subnet-insert

Indicates, when enabled, that the system should set the edns0 client subnet option to the source address for queries that do not already contain a client subnet option. Also specifies that the system should remove the client subnet option from responses to clients that did not send a client subnet option in their most recent query.

enable-cache

Indicates whether the system caches DNS responses. The default value is no.

enable-dnssec

Indicates whether to perform DNS Security Extension (DNSSEC) operations on the DNS packet, for example, respond to DNSKEY queries; add RRSIGs to response.

enable-dns-express

Indicates whether the dns-express service is enabled. The service handles zone transfers from the primary DNS server.

enable-dns-firewall

Indicates whether DNS firewall capability is enabled. The default value is no.

enable-gtm

Indicates whether the Global Traffic Manager handles DNS resolution for DNS queries and responses that contain Wide IP names. The default value is yes.

enable-hardware-query-validation

On supported platforms, indicates whether the hardware will accelerate query validation. The default value is no.

enable-hardware-response-cache

On supported platforms, indicates whether the hardware will cache responses. The default value is no.

enable-logging

Indicates whether to enable high speed logging for DNS queries and responses or not. Default value is no. When it is set to yes, a DNS profile must be configured with a log-profile.

enable-rapid-response

On supported platforms, indicates whether to allow queries to be answered by Rapid Response. The default value is no. When enabled, if the query name matches a GTM Wide IP name and GTM is enabled on this profile, the DNS query will bypass Rapid Response.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

log-profile

Specifies the DNS logging profile used to configure what events get logged and their message format.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the profile resides.

process-rd
Indicates whether to process client-side DNS packets with Recursion Desired set in the header. The default value is yes. If set to no, processing of the packet will be subject to the unhandled-query-action option.

process-xfr
Indicates whether the system answers zone transfer requests for a DNS zone created on the system. The default value is no. The enable-dns-express and process-xfr settings affect how the system responds to zone transfer requests.

rapid-response-last-action
Specifies what action to take when Rapid Response is enabled and the incoming query has not matched a DNS-Express Zone. Default is drop. Option allow sends non-matching queries up the regular packet processing path. All other options result in a response returned immediately to the client: truncate (truncate), nxdomain (non-existent name), noerror (no data), refuse (REFUSED return code).

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

to-folder
dns profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

unhandled-query-action
Specifies the action to take when a query does not match a Wide IP or a DNS Express Zone. The default value is allow.

use-local-bind
Indicates whether non-GTM and non-dns-express requests should be forwarded to the local BIND.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-05-26 ltm profile dns(1)

ltm profile fasthttp

NAME
fasthttp - Configures a Fast HTTP profile.

MODULE
ltm profile

SYNTAX
Modify the fasthttp component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY
create fasthttp [name]
modify fasthttp [name]
options:
app-service [[string] | none]
client-close-timeout [integer]
connpool-idle-timeout-override [integer]
connpool-max-reuse [integer]
connpool-max-size [integer]
connpool-min-size [integer]
connpool-replenish [disabled | enabled]
connpool-step [integer]
defaults-from [[name] | none]
description [string]
force-http-10-response [disabled | enabled]
hardware-syn-cookie [disabled | enabled]
header-insert [none | [string]]

http-11-close-workarounds [disabled | enabled]
idle-timeout [integer]
insert-xforwarded-for [disabled | enabled]
layer-7 [disabled | enabled]
max-header-size [integer]
max-requests [integer]
mss-override [integer]
receive-window-size [65535 - 2³¹ bytes for window scale enabling]
reset-on-timeout [disabled | enabled]
server-close-timeout [integer]
server-sack [disabled | enabled]
server-timestamp [disabled | enabled]
unclean-shutdown [disabled | enabled]

mv fasthttp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]
options:
to-folder

edit fasthttp [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

reset-stats fasthttp
reset-stats fasthttp [[[name] | [glob] | [regex]] ...]

DISPLAY
list fasthttp
list fasthttp [[[name] | [glob] | [regex]] ...]
show running-config fasthttp
show running-config fasthttp [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

show fasthttp
show fasthttp [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete fasthttp [name]

DESCRIPTION

You can use this component to create, modify, display, or delete a Fast HTTP profile. This profile provides the ability to accelerate certain HTTP connections such as banner ads.

EXAMPLES

```
create fasthttp my_fast_http_profile defaults-from fasthttp
```

Creates a Fast HTTP profile named `my_fast_http_profile` that inherits its settings from the system default Fast HTTP profile.

```
mv fasthttp /Common/my_fasthttp_profile to-folder /Common/my_folder
```

Moves a custom fasthttp profile named `my_fasthttp_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

```
show fasthttp
```

Displays fasthttp profile statistics in the system default units.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`client-close-timeout`

Specifies the number of seconds after which the system closes a client connection, when the system either receives a client FIN packet or sends a FIN packet. This option overrides the `idle-timeout` option. The default value is 5.

`server-sack`

Specifies whether to support server sack option in cookie response by default. The default value is disabled.

`server-timestamp`

Specifies whether to support server timestamp option in cookie response by default. The default value is disabled.

`receive-window-size`

Specifies the window size to use, minimum and default to 65535 bytes, the maximum is 2³¹ for window

scale enabling.

`connpool-idle-timeout-override`

Specifies the number of seconds after which a server-side connection in a OneConnect(tm) pool is eligible for deletion, when the connection has no traffic. This option overrides the `idle-timeout` option. The default value is 0 (zero) seconds, which disables the override setting.

`connpool-max-reuse`

Specifies the maximum number of times that the system can re-use a current connection. The default value is 0 (zero).

`connpool-max-size`

Specifies the maximum number of connections to a load balancing pool. A value of 0 (zero) specifies that a pool can accept an unlimited number of connections. The default value is 2048.

`connpool-min-size`

Specifies the minimum number of connections to a load balancing pool. The default value of 0 (zero) specifies that there is no minimum.

`connpool-replenish`

When enabled, the system replenishes the number of connections to a load balancing pool to the number of connections that existed when the server closed the connection to the pool. The default value is enabled.

When disabled, the system replenishes the connection that was closed by the server, only when there are fewer connections to the pool than the number of connections set in the `connpool-min-size` option.

`connpool-step`

Specifies the increment at which the system makes additional connections available, when all available connections are in use. The default value is 4.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is `fasthttp`.

`description`

User defined description.

`force-http10-response`

Specifies whether to rewrite the HTTP version in the status line of the server to HTTP 1.0 to discourage the client from pipelining or chunking data. The default value is disabled.

`glob` Displays the items that match the `glob` expression. See `help glob` for a description of `glob` expression syntax.

`hardware-syn-cookie`

This option is deprecated in version 13.0.0 because it is not a fully functioning feature. Specifies whether or not to use hardware SYN Cookie when cross system limit. The default value is enabled

`header-insert`

Specifies a string that the system inserts as a header in an HTTP request. If the header exists already, the system does not replace it. The default value is none.

`http11-close-workarounds`

Enables or disables HTTP 1.1 close workarounds. The default value is disabled.

`idle-timeout`

Specifies the number of seconds after which a connection is eligible for deletion, when the connection has no traffic. The default value is 300 seconds.

`insert-xforwarded-for`

Specifies whether the system inserts the XForwarded For header in an HTTP request with the client IP address, to use with connection pooling.

The options are:

`disabled`

Specifies that the system does not insert the XForwarded For header.

`enabled`

Specifies that the system inserts the XForwarded For header with the client IP address.

`layer7`

When enabled, the system parses HTTP data in the stream. Disable this option if you want to use the performance HTTP profile to shield against denial-of-service attacks against non-HTTP protocols. The default value is enabled.

`max-header-size`

Specifies the maximum amount of HTTP header data that the system buffers before making a load balancing decision. The default value is 32768.

`max-requests`

Specifies the maximum number of requests that the system can receive on a client connection, before the system closes the connection. The default value of 0 specifies that requests are not limited.

`mss-override`

Specifies a maximum segment size (MSS) override for server connections. The default value is 0 (zero), which corresponds to an MSS of 1460. You can specify any integer between 536 and 1460.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reset-on-timeout
When enabled, the system sends a TCP RESET packet when a connection times out, and deletes the connection. The default value is enabled.

server-close-timeout
Specifies the number of seconds after which the system closes a client connection, when the system either receives a client FIN packet or sends a FIN packet. This option overrides the value of the idle-timeout option. The default value is 5.

to-folder
fasthttp profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

unclean-shutdown
Specifies how the system handles closing a connection. The options are:

disabled
Prevents an unclean shutdown of a client connection. This is the default value.

enabled
Specifies to permit an unclean shutdown of a client connection.

fast Specifies that the system sends a RESET packet to close the connection only if the client attempts to send further data after the response has completed.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2016-05-10 ltm profile fasthttp(1)

ltm profile fastl4

NAME

fastl4 - Configures a Fast Layer 4 profile.

MODULE

ltm profile

SYNTAX

Configure the fastl4 component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create fastl4 [name]

modify fastl4 [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

hardware-syn-cookie [disabled | enabled]

idle-timeout [immediate | indefinite | [integer]]

ip-tos-to-client [[integer] | pass-through]

ip-tos-to-server [[integer] | pass-through]

keep-alive-interval [integer]

ip-df-mode [preserve | set | clear]

ip-ttl-mode [proxy | preserve | decrement | set]

ip-ttl-value [integer]

link-qos-to-client [[integer] | pass-through]

link-qos-to-server [[integer] | pass-through]

priority-to-client [[integer] | pass-through]

priority-to-server [[integer] | pass-through]

loose-close [disabled | enabled]
loose-initialization [disabled | enabled]
mss-override [integer]
pva-acceleration [full | none | partial | dedicated]
pva-dynamic-client-packets [integer]
pva-dynamic-server-packets [integer]
pva-offload-dynamic [enabled | disabled]
pva-offload-state [embryonic | establish]
pva-offload-dynamic-priority [enable | disable]
pva-offload-initial-priority [low | medium | high]
pva-flow-aging [enabled | disabled]
pva-flow-evict [enabled | disabled]
tcp-pva-when-to-offload [embryonic | establish]
tcp-pva-offload-direction [bidirectional | client-to-server-only | server-to-client-only]
other-pva-when-to-offload [after-packets-per-direction | after-packets-both-direction]
other-pva-offload-direction [bidirectional | client-to-server-only | server-to-client-only]
other-pva-clientpkts-threshold [integer]
other-pva-serverpkts-threshold [integer]
reassemble-fragments [disabled | enabled]
reset-on-timeout [disabled | enabled]
rtt-from-client [disabled | enabled]
rtt-from-server [disabled | enabled]
server-sack [disabled | enabled]
server-timestamp [disabled | enabled]
receive-window-size [65535 - 2³¹ bytes for window scale enabling]
software-syn-cookie [disabled | enabled]
syn-cookie-dsr-flow-reset-by [bigip | client | none]
syn-cookie-enable [disabled | enabled]
syn-cookie-mss [integer]
syn-cookie-whitelist [disabled | enabled]
tcp-close-timeout [immediate | indefinite | [integer]]
tcp-generate-is [disabled | enabled]
tcp-handshake-timeout [immediate | indefinite | [integer]]
tcp-strip-sack [disabled | enabled]
tcp-timestamp-mode [preserve | rewrite | strip]
tcp-time-wait-timeout [integer]
tcp-wscale-mode [preserve | rewrite | strip]
late-binding [enabled | disabled]
explicit-flow-migration [enabled | disabled]
client-timeout [integer]
timeout-recovery [disconnect | fallback]

mv fastl4 [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]
options:
to-folder

edit fastl4 [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

reset-stats fastl4
reset-stats fastl4 [[[name] | [glob] | [regex]] ...]

DISPLAY
list fastl4
list fastl4 [[[name] | [glob] | [regex]] ...]
show running-config fastl4
show running-config fastl4
[[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

show fastl4
show fastl4 [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete fastL4 [name]

DESCRIPTION

You can use this component to create, modify, display, or delete a Fast Layer 4 profile. The Fast L4 profile is the default profile that the system uses when you create a basic configuration for non-UDP (User Datagram Protocol) traffic.

Any changes you make to an active Fast L4 profile (one that is in use by a virtual server) take effect after the value of the idle-timeout option has passed. That means new connections are affected by the profile change immediately. However, for the new values to take effect, old connections need to be either aged out or closed.

EXAMPLES

create fastl4 my_fastl4_profile defaults-from fastl4

Creates a custom Fast Layer 4 profile named `my_fastl4_profile` that inherits its settings from the system default Fast L4 profile.

```
mv fastl4 /Common/my_fastl4_profile to-folder /Common/my_folder
```

Moves a custom fastl4 profile named `my_fastl4_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

Please refer to the mv manual page for examples on how to use the mv command.

```
show fastl4
```

Displays statistics for all Fast Layer 4 profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is fastl4.

description

User defined description.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

hardware-syn-cookie

This option is deprecated in version 13.0.0 and is replaced by `syn-cookie-enable`. Enables or disables hardware SYN cookie support when PVA10 is present on the system. The default value is disabled.

Note that when you set the hardware-syn-cookie option to enabled, you may also want to set the following bigdb database variables using the db component, based on your requirements:

```
pva.SynCookies.Full.ConnectionThreshold (default: 500000)
pva.SynCookies.Assist.ConnectionThreshold (default: 500000)
pva.SynCookies.ClientWindow (default: 0)
```

idle-timeout

Specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 300 seconds. You can also specify immediate or indefinite.

When you specify an idle-timeout for the Fast L4 profile, for the profile to work properly, the value needs to be greater than the bigdb database variable `Pva.Scrub_time_in_msec`.

ip-tos-to-client

Specifies an IP Type of Service (ToS) number for the client-side. This option specifies the ToS level that the traffic management system assigns to IP packets when sending them to clients. The default value is 65535, which indicates, do not modify.

ip-tos-to-server

Specifies an IP ToS number for the server side. This option specifies the ToS level that the traffic management system assigns to IP packets when sending them to servers. The default value is 65535, which indicates, do not modify.

keep-alive-interval

Specifies the keep-alive probe interval, in seconds. The default value is disabled (0 seconds).

ip-df-mode

Describe the Don't Fragment (DF) bit setting in the IP Header of the outgoing TCP packet. The available settings are: `Pmtu`: Set the outgoing IP Header DF bit based on IP pmtu setting (`tm.pathmtudiscovery`). `Preserve`: Set the outgoing Packet's IP Header DF bit to be same as incoming IP Header DF bit. `Set`: Set the outgoing packet's IP Header DF bit. `Clear`: Clear the outgoing packet's IP Header DF bit. The default setting is Preserve.

ip-ttl-mode

Describe the outgoing TCP packet's IP Header TTL mode. The available Modes are: `Proxy`: Set the outgoing IP Header TTL value to 255/64 for ipv4/ipv6 respectively. `Preserve`: Set the outgoing IP Header TTL value to be same as the incoming IP Header TTL value. `Decrement`: Set the outgoing IP Header TTL value to be one less than the incoming TTL value. `Set`: Set the outgoing IP Header TTL value to a specific value (as specified by `ip-ttl-v[4|6]`). The default mode is Decrement.

ip-ttl-v4

Specify the outgoing packet's IP Header TTL value for IPv4 traffic. Maximum TTL value that can be specified is 255. The default is 255.

ip-ttl-v6

Specify the outgoing packet's IP Header TTL value for IPv6 traffic. Maximum TTL value that can be specified is 255. The default is 64.

link-qos-to-client

Specifies a Link Quality of Service (QoS) (VLAN priority) number for the client side. This option specifies the QoS level that the system assigns to packets when sending them to clients. The default value is 65535, which indicates, do not modify.

link-qos-to-server

Specifies a Link QoS (VLAN priority) number for the server side. This option specifies the QoS level that the system assigns to packets when sending them to servers. The default value is 65535, which indicates, do not modify.

priority-to-client

Specifies internal packet priority for the client side. This option specifies the internal packet priority that the system assigns to packets when sending them to clients. The default value is 65535, which indicates, do not modify.

priority-to-server

Specifies internal packet priority for the server side. This option specifies the internal packet priority that the system assigns to packets when sending them to servers. The default value is 65535, which indicates, do not modify.

loose-close

Specifies that the system closes a loosely-initiated connection when the system receives the first FIN packet from either the client or the server. The default value is disabled.

loose-initialization

Specifies that the system initializes a connection when it receives any Transmission Control Protocol (TCP) packet, rather than requiring a SYN packet for connection initiation. The default value is disabled.

mss-override

Specifies a maximum segment size (MSS) override for server connections. Note that this is also the MSS advertised to a client when a client first connects.

The default value is 0 (zero), which disables this option. You can specify an integer from 256 to 9162.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

pva-acceleration

Specifies the Packet Velocity(r) ASIC acceleration policy. The default value is full. In 12.1, dedicated mode is the new low latency policy which replaces guaranteed. The full and partial mode has same effect for ePVA platforms.

pva-dynamic-client-packets

Specifies the number of client packets before dynamic ePVA hardware re-offloading occurs. The valid value is 0–10. The default value is 2.

pva-dynamic-server-packets

Specifies the number of server packets before dynamic ePVA hardware re-offloading occurs. The valid value is 0–10. The default value is 2.

pva-offload-dynamic

Specifies whether PVA flow dynamic offloading is enabled or not. The default is enabled.

For a flow or flow(s) in a connection to be offloaded to ePVA hardware, both the client (pva-dynamic-client-packets) and server (pva-dynamic-server-packets) flow packets setting need to be satisfied. If only one direction packets need to be taken into consideration, the other direction packets should set to zero.

pva-offload-initial-priority

Specifies the initial epva offload priority of a flow. Priority can be low, medium or high. The default value is medium

pva-offload-dynamic-priority

Specifies if dynamic adjustment of epva offload flow priority is turned on or not. Default value is disabled.

pva-offload-state

This option is deprecated in version 14.1.0 and is replaced by tcp-pva-when-to-offload and other-pva-when-to-offload. Specifies at what stage the ePVA performs hardware offload. The default value is embryonic and implies at TCP CSYN or the first client UDP packet. establish implies TCP 3WAY handshaking or UDP CS round trip are confirmed.

pva-flow-aging

Specifies if automatic aging from ePVA flow cache upon inactive and idle for a period, default to enabled.

pva-flow-evict

Specifies if this flow can be evicted upon hash collision with a new flow learn snoop request, defaults to enabled.

tcp-pva-when-to-offload

Specifies at what stage the ePVA performs hardware offload for TCP traffic. The default value is embryonic and implies at TCP SYN packet. establish implies TCP 3WAY handshaking.

tcp-pva-offload-direction

For tcp protocol traffic only, specifies which side of the traffic can ePVA perform hardware offload for. The default value is bidirectional which implies both side is permitted to offload if threshold exceeds.

client-to-server-only implies only the traffic from client to server is allowed to be offloaded. Even if the traffic from server to client exceeds the threshold, it will not be offloaded. Vice versa, server-to-client-only implies only the traffic from server to client is allowed to be offloaded.

other-pva-whento-offload

Specifies when the ePVA performs hardware offload for stateless protocol traffic. The default value is after-packets-per-direction and implies the client and server traffic is offloaded independently after exceeding their own thresholds. after-packets-both-direction implies both client and server traffic thresholds need to be exceeded, then can both sides get offloaded.

other-pva-offload-direction

For stateless protocol traffic only, specifies which side of the traffic can ePVA perform hardware offload for. The default value is bidirectional which implies both side is permitted to offload if threshold exceeds. client-to-server-only implies only the traffic from client to server is allowed to be offloaded. Even if the traffic from server to client exceeds the threshold, it will not be offloaded. Vice versa, server-to-client-only implies only the traffic from server to client is allowed to be offloaded.

other-pva-clientpkts-threshold

Specifies the number of client packets before ePVA hardware offloading occurs for stateless protocol traffic. The valid value is 0~255. The default value is 2.

other-pva-serverpkts-threshold

Specifies the number of server packets before ePVA hardware offloading occurs for stateless protocol traffic. The valid value is 0~255. The default value is 1.

reassemble-fragments

Specifies whether to reassemble fragments. The default value is disabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reset-on-timeout

Specifies whether you want to reset connections on timeout. The default value is enabled.

rtt-from-client

Enables or disables the TCP timestamp options to measure the round trip time to the client. The default value is disabled.

rtt-from-server

Enables or disables the TCP timestamp options to measure the round trip time to the server. The default value is disabled.

server-sack

Specifies whether to support server sack option in cookie response by default. The default value is disabled.

server-timestamp

Specifies whether to support server timestamp option in cookie response by default. The default value is disabled.

receive-window-size

Specifies the window size to use, minimum and default to 65535 bytes, the maximum is 2^{31} for window scale enabling.

software-syn-cookie

This option is deprecated in version 13.0.0 and is replaced by syn-cookie-enable. Enables or disables software SYN cookie support when PVA10 is not present on the system. The default value is disabled.

syn-cookie-dsr-flow-reset-by

Specifies how TCP SYN Flood is handled when syn-cookie-whitelist is enabled and the attack is detected in Direct Server Return(DSR) mode. The default value is none, which is backward-compatible with syn-cookie-whitelist feature in non-DSR mode.

syn-cookie-enable

Enables syn-cookies capability on this virtual server. For the details on the threshold at which syn-cookies are triggered please see default-vs-syn-challenge-threshold and global-syn-challenge-threshold or the tcp-half-open vector in the DoS profile. The default is enabled.

syn-cookie-mss

Specifies a maximum segment size (MSS) for server connections when SYN Cookie is enabled. Note that this is also the MSS advertised to a client when a client first connects.

The default value is 0 (zero), which disables this option. You can specify an integer from 256 to 9162.

syn-cookie-whitelist

Specifies whether or not to use a SYN Cookie WhiteList when doing software SYN Cookies. This means not doing a SYN Cookie for the same src IP address if it has been done already in the previous tm.flowstate.timeout (30) seconds. The default value is disabled.

tcp-close-timeout

Specifies a TCP close timeout in seconds. You can also specify immediate or indefinite. The default value is 5 seconds.

tcp-generate-isn

Specifies whether you want to generate TCP sequence numbers on all SYNs that conform with RFC1948, and allow timestamp recycling. The default value is disabled.

tcp-handshake-timeout

Specifies a TCP handshake timeout in seconds. You can also specify immediate or indefinite. The default value is 5 seconds.

tcp-time-wait-timeout

Specifies a TCP time_wait timeout in milliseconds. The default value is 0 milliseconds.

tcp-strip-sack

Specifies whether you want to block the TCP SackOK option from passing to the server on an initiating SYN. The default value is disabled.

tcp-timestamp-mode

Specifies how you want to handle the TCP timestamp. The default value is preserve.

tcp-wscale-mode

Specifies how you want to handle the TCP window scale. The default value is preserve.

late-binding

Specifies whether to enable or disable intelligent selection of a back-end server pool. The default value is disabled. With this option enabled, an iRule can read a Layer 7 (FIX) packet to select a server pool, and then can send the FIX stream down to the ePVA. The ePVA then manages the FIX stream at a low latency, for as long as the stream persists. To keep the latency low, the BIG-IP software does not examine any more Layer-7 data in that FIX stream.

If you enable this option, you also need a FIX profile in the Performance FastL4 Virtual Server configuration.

explicit-flow-migration

Specifies whether to have the iRule code determine exactly when the FIX stream drops down to the ePVA hardware. The default value is disabled.

The explicit flow migration state indicates whether connections are automatically migrated into the ePVA hardware (disabled), or the iRule must explicitly migrate them with the BIGTCP::release_flow command (enabled).

client-timeout

Specifies late binding client timeout in seconds. This is the number of seconds allowed for a client to transmit enough data to select a server pool. If this timeout expires, the timeout-recovery option dictates whether to drop the connection or fallback to the normal FastL4 load-balancing method to pick a server pool. The default timeout is 30 seconds.

timeout-recovery

Specifies late binding timeout recovery mode. This is the action to take when late binding timeout occurs on a connection. This could be disconnect if only the L7 iRule actions are acceptable to pick a server or fallback if the normal FastL4 load-balancing methods are acceptable to pick a server. The default action is to disconnect.

to-folder

fastl4 profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2020-02-12 ltm profile fastl4(1)

Itm profile fix

NAME

fix - Configures an Financial Information eXchange Protocol (FIX) profile.

MODULE

ltm profile

SYNTAX

Configure the fix component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create fix [name]

modify fix [name]
options:
app-service [[string] | none]
defaults-from [[name] | none]
description [string]
error-action [drop_connection | dont_forward]
full-logon-parsing [true | false]
message-log-publisher [publisher]
quick-parsing [true | false]
statistics-sample-interval [integer]
report-log-publisher [publisher]
response-parsing [true | false]
sender-tag-class {[sender-id] [class name]}...}

edit fix [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

reset-stats fix
reset-stats fix [[[name] | [glob] | [regex]] ...]

DISPLAY
list fix
list fix [[[name] | [glob] | [regex]] ...]
show running-config fix
show running-config fix [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

show fix
show fix [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE
delete fix [name]

DESCRIPTION

You can use the fix component to manage an Financial Information eXchange Protocol profile.

EXAMPLES

```
create fix my_fix defaults-from fix
```

Creates an financial information exchange protocol profile named my_fix using the system defaults.

```
create fix my_fix { }
```

Creates an financial information exchange protocol profile named my_fix.

app-service

Specifies the name of the application service to which this object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is none.

description

User defined description.

error-action

Specifies the error handling method.

full-logon-parsing

Enable or disable logon message is always fully parsed.

message-log-publisher

Specifies the publisher for message logging.

quick-parsing

Enable or disable quick parsing which parses the basic standard fields and validates message length and checksum.

statistics-sample-interval

Specifies the sample interval in seconds of the message rate.

response-parsing

Enable or disable response parsing which parses the messages from FIX server.

report-log-publisher

Specifies the publisher for error message and status report.

partition

Specifies the administrative partition within which the profile resides.

sender-tag-class

Specifies the tag substitution map between sender id and tag substitution data group.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh, ltm profile fix

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013. All rights reserved.

BIG-IP 2014-10-20 ltm profile fix(1)

ltm profile ftp

NAME

ftp - Configures an FTP profile.

MODULE

ltm profile

SYNTAX

Configure the ftp component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create ftp [name]

modify ftp [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

port [name]

allow-ftps [disabled | enabled]

ftps-mode [disallow | allow | require]

enforce-tls-session-reuse [disabled | enabled]

allow-active-mode [disabled | enabled]

security [disabled | enabled]

translate-extended [disabled | enabled]

inherit-parent-profile [disabled | enabled]

log-publisher [log publisher name | none]

log-profile [log profile name | none]

mv ftp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

edit ftp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ftp

list ftp [[[name] | [glob] | [regex]] ...]

show running-config ftp

show running-config ftp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete ftp [name]

DESCRIPTION

Use this command to create, modify, display, or delete an FTP profile with which you can manage FTP traffic.

EXAMPLES

create ftp my_ftp_profile defaults-from ftp

Creates a custom FTP profile named my_ftp_profile that inherits its settings from the system default FTP

profile.

list ftp

Displays the properties of all FTP profiles.

```
mv ftp /Common/my_ftp_profile to-folder /Common/my_folder
```

Moves a custom ftp profile named my_ftp_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is ftp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

port Specifies a service for the data channel port used for this FTP profile. The default port is ftp-data.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

allow-ftps

Deprecated since v14.0.0. Use ftps-mode instead. Allow explicit FTPS negotiation. The default value is disabled.

ftps-mode

Specifies the policy for explicit FTPS negotiation on FTP command channel, including disallow, allow and require. The default is disallow.

Mode disabled

Detects explicit FTPS negotiation, and AUTH TLS command will be dropped.

Mode

Bypasses explicit FTPS negotiation through the BIGIP.

Mode allow

Allows FTPS negotiation, and will intercept TLS if SSL forward proxy is configured on the virtual server.

Mode require

Requires FTPS negotiation to protect the FTP data transfers.

enforce-tls-session-reuse

Enforce data connection to reuse TLS session. The default value is disabled.

allow-active-mode

Allow FTP active transfer mode. The default value is enabled.

security

Enables or disables secure FTP traffic for the BIG-IP(r) Application Security Manager. You can set the security option only if the system is licensed for the BIG-IP Application Security Manager. The default value is disabled.

to-folder

ftp profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

translate-extended

This option is enabled by default, and thus, automatically translates RFC2428 extended requests EPSV and EPRT to PASV and PORT when communicating with IPv4 servers.

inherit-parent-profile

Enables the FTP data channel to inherit the TCP profile used by the control channel. If disabled, the data channel uses FastL4 (BigProto) only.

log-publisher

Specify the name of the log publisher which logs translation events. See help sys log-config for more details on the logging sub-system. Use the "sys log-config publisher" component to set up a log

publisher.

log-profile

Specify the name of the ALG log profile which controls the logging of ALG . See help ltm alg-log-profile for more details on the logging profile sub-system. Use the "ltm alg-log-profile profile" component to set up a ALG log profile.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2018-04-02 ltm profile ftp(1)

ltm profile georedundancy

NAME

georedundancy - Configures a Geo-Redundancy profile.

MODULE

ltm profile

SYNTAX

Configure the georedundancy component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create georedundancy [name]

modify georedundancy [name]

options:

defaults-from [[name] | none]

description [string]

group-id [string]

local-site-id [string]

prefix [[string] | none]

metadata-refresh-interval-ms [integer]

message-send-max-retries [integer]

message-timeout-ms [integer]

read-broker-list [string]

remote-site-id [string]

transport-name [string]

write-broker-list [string]

edit georedundancy [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv georedundancy [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats georedundancy

reset-stats georedundancy [[[name] | [glob] | [regex]] ...]

DISPLAY

list georedundancy

list georedundancy [[[name] | [glob] | [regex]] ...]

show running-config georedundancy

show running-config georedundancy [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show georedundancy

show georedundancy [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete georedundancy [name]

DESCRIPTION

You can use the georedundancy component to manage a Geo-Redundancy profile.

EXAMPLES

```
create georedundancy my_georedundancy_profile defaults-from georedundancy
```

Creates a Geo-Redundancy profile named my_georedundancy_profile using the system defaults.

```
create georedundancy my_georedundancy_profile { local-site-id local_site }
```

Creates a Geo-Redundancy profile named my_georedundancy_profile with a name of the local site identification set to local_site.

```
mv georedundancy /Common/my_georedundancy_profile to-folder /Common/my_folder
```

Moves a custom georedundancy profile named my_georedundancy_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

local-site-id

Defines the name to identify local site . The default value is local_site.

remote-site-ids

Defines the list of comma separated names corresponding to remote sites. The default value is remote_site.

read-broker-list

Specifies the list of local site broker's IP addresses and ports for the local TMMs to read the data from. Example: 127.20.1.254:9092, 127.20.2.254:9092.

write-broker-list

Specifies the list of local site broker's IP addresses and ports for the local TMMs to write the data to. Example: 127.20.1.254:9092, 127.20.2.254:9092.

transport-name

Defines the transport profile that Geo-Redundancy half proxy uses to establish connection to brokers. The default value is tcp.

metadata-refresh-interval-ms

Indicates the Topic metadata refresh interval in milliseconds. The metadata is automatically refreshed on error and connect. Use -1 to disable the intervalled refresh. The default value is 300K.

group-id

Defines the client group identifier string. All clients sharing the same group-id belong to the same group.

message-send-max-retries

Specifies the maximum number of attempts that will be made to reload the Diameter session DB from the remote site during startup.

message-timeout-ms

Specifies how often the system will attempt to reload the Diameter session DB from the remote site during startup.

prefix

Indicates a list of session-db key prefixes and for each prefix a session-db global callback gets registered. A prefix when applied to any logical table within the iRule session-db storage would enable replication of the same. These prefix values will only be applied to iRule session bins. The default value is "".

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2016-2017. All rights reserved.

BIG-IP 2019-08-30 ltm profile georedundancy(1)

ltm profile gtp

NAME

gtp - Configures a GTP profile.

MODULE

ltm profile

SYNTAX

Configure the gtp component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create gtp [name]

modify gtp [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

ingress-max [integer]

edit gtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv gtp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats gtp

reset-stats gtp [[[name] | [glob] | [regex]] ...]

DISPLAY

list gtp

list gtp [[[name] | [glob] | [regex]] ...]

show running-config gtp

show running-config gtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show gtp

show gtp [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete gtp [name]

DESCRIPTION

You can use the gtp component to manage a GTP profile.

EXAMPLES

```
create gtp my_gtp_profile defaults-from gtp
```

Creates a GTP profile named my_gtp_profile using the system defaults.

```
create gtp my_gtp_profile { ingress-max 1000 }
```

Creates a GTP profile named my_gtp_profile that specifies the maximum number of messages that can be held in the ingress queue is 1000.

```
mv gtp /Common/my_gtp_profile to-folder /Common/my_folder
```

Moves a custom gtp profile named my_gtp_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is gtp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ingress-max

Specifies the maximum number of messages that can be held in ingress queue. If it is 0, then it is unlimited. The default value is 0.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2015. All rights reserved.

BIG-IP 2015-06-01 ltm profile gtp(1)

ltm profile html

NAME

html - Configures an HTML profile.

MODULE

ltm profile

SYNTAX

Configure the html component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create html [name]

modify html [name]

options:

defaults-from [[name] | none]

content-detection [disabled | enabled]

content-selection

[add | delete | replace-all-with] {

[content-type] ...

}

content-selection none

rules

[add | delete | replace-all-with] {

[html-rule] ...

}

rules none

mv html [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats html

reset-stats html [[[name] | [glob] | [regex]] ...]

DISPLAY

list html

list html [[[name] | [glob] | [regex]] ...]

show running-config html

show running-config html [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show html

show html [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete html [name]

DESCRIPTION

Use this command to create, modify, display, or delete an HTML profile with which you can manage HTML traffic.

EXAMPLES

```
create html my_html_profile defaults-from html
```

Creates a custom HTML profile named `my_html_profile` that inherits its settings from the system default HTML profile.

```
list html
```

Displays the properties of all HTML profiles.

```
mv html /Common/my_html_profile to-folder /Common/my_folder
```

Moves a custom HTML profile named `my_html_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

Please refer to the `mv` manual page for examples on how to use the `mv` command.

OPTIONS

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is `html`.

`description`

User defined description.

`glob` Displays the items that match the `glob` expression. See `help glob` for a description of `glob` expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`content-detection`

Scans initial HTTP payload to look for HTML signatures and enables HTML profile if HTML-like patterns are detected.

`content-selection`

Matches content-type from response header against a list of content-types and enables HTML profile if a match is found.

`rules`

Specifies a list of HTML (content rewrite) rules, separated by spaces, that are used for parsing and patching HTML.

`to-folder`

HTML profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

SEE ALSO

`create`, `delete`, `glob`, `list`, `itm virtual`, `modify`, `mv`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 itm profile html(1)

Itm profile http-compression

NAME

`http-compression` - Configures an HTTP Compression profile.

MODULE

`itm profile`

SYNTAX

Configure the `http-compression` component within the `itm profile` module using the syntax shown in the following sections.

CREATE/MODIFY

```
create http-compression [name]
```

```
modify http-compression [name]
```

options:

```
allow-http-10 [disabled | enabled]
```

```

app-service [[string] | none]
browser-workarounds [disabled | enabled]
buffer-size [integer]
cpu-saver [disabled | enabled]
cpu-saver-high [integer]
cpu-saver-low [integer]
content-type-exclude
  [add | delete | replace-all-with] {
[content type] ...
  }
content-type-exclude none
content-type-include
  [add | delete | replace-all-with] {
[content type] ...
  }
content-type-include none
defaults-from [ [name] | none]
description [string]
gzip-level [integer]
gzip-memory-level [integer, in bytes]
gzip-window-size [integer]
keep-accept-encoding [disabled | enabled]
method-prefer [deflate | gzip]
min-size [integer]
selective [disabled | enabled]
uri-exclude
  [add | delete | replace-all-with] {
[URI] ...
  }
uri-exclude none
uri-include
  [add | delete | replace-all-with] {
[URI] ...
  }
uri-include none
vary-header [disabled | enabled]

edit http-compression [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

mv http-compression [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
options:
  to-folder

reset-stats http-compression
reset-stats http-compression [ [ [name] | [glob] | [regex] ] ... ]

DISPLAY
list http-compression
list http-compression [ [ [name] | [glob] | [regex] ] ... ]
show running-config http-compression
show running-config http-compression [ [ [name] | [glob] | [regex] ]
... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

show http-compression
show http-compression [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global

DELETE
delete http-compression [name]

```

DESCRIPTION

You can use the http-compression component to create, modify, display, or delete an HTTP Compression profile.

The BIG-IP(r) system installation includes the following default HTTP Compression-type profiles:

• http-compression

• wan-optimized-compression

The default HTTP Compression profile contains values for properties related to managing compression settings.

You can create a new HTTP Compression-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

```
create http-compression my_hc_profile defaults-from http-compression
```

Creates a custom HTTP Compression profile named `my_hc_profile` that inherits its settings from the system default HTTP Compression profile.

```
mv http-compression /Common/my_httpcompression_profile to-folder /Common/my_folder
```

Moves a custom `http-compression` profile named `my_httpcompression_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`allow-http10`

Enables or disables compression of HTTP/1.0 server responses. The default value is disabled.

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`browser-workarounds`

Deprecated since v13.1.0. Enables or disables browser workarounds. The default value is disabled.

• If the client browser is Netscape Navigator(r) version 4.0x, compression is turned off. Netscape advertises that the browser can handle compression gracefully. In this case, F5 Networks disables compression entirely for that class of browser.

• If the client browser is Netscape Navigator version 4.x (4.10 and later) and the server response Content-Type is not either `text/html` or `text/plain` compression is turned off. This class of Netscape browsers can handle plain text and HTML just fine, but there are known issues with other types of content.

• If the client browser is Microsoft(r) Internet Explorer (any version), the server response Content-Type is either `text/css` or `application/x-javascript`, and the client connection is over SSL, compression is turned off. The Microsoft article ID for this problem is 825057.

• If the client browser is Microsoft Internet Explorer (any version), the server response Content-Type is either `text/css` or `application/x-javascript`, and the server sets the header `Cache-Control` to `no-cache`, compression is turned off. The Microsoft article ID for this problem is 327286.

`buffer-size`

Specifies the maximum number of uncompressed bytes that the system buffers before determining whether to compress the response. Useful when the headers of a server response do not specify the length of the response content. The default value is 4096.

`content-type-exclude`

Specifies a string list of HTTP Content-Type responses that you do not want the system to compress. The default value is none.

`content-type-include`

Specifies a string list of HTTP Content-Type responses that you want the system to compress. The default value is `{ text/ application/ (xml|x-javascript) }`.

`cpu-saver`

Enables or disables the CPU saver feature. When the CPU saver is enabled, the system monitors the percent of CPU usage and adjusts compression rates automatically when the CPU usage reaches the percentage defined in the `compress-cpu-saver-low` and `compress-cpu-saver-high` options. The default value is enabled.

`cpu-saver-high`

Specifies the percent of CPU usage at which the system starts automatically decreasing the amount of content being compressed, as well as the amount of compression that the system is applying. The default value is 90.

`cpu-saver-low`

Specifies the percent of CPU usage at which the system resumes content compression at the user-defined rates. The default value is 75.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is `httpcompression`.

`description`

User defined description.

`gzip-level`

Specifies a value that determines the amount of memory that the system uses when compressing a server response. The default value is 1.

`gzip-memory-level`

Specifies the amount of memory (in kilobytes) that the system uses when compressing a server response. The system rounds the value up to the nearest power of two. The default value is 8. The maximum value is 256.

`gzip-window-size`

Specifies the number of kilobytes in the window size that the system uses when compressing a server response. The system rounds the value up to the nearest power of two. The default value is 16k. The maximum value is 128k.

`keep-accept-encoding`

Specifies where data compression is performed. When enabled, the target server, rather than the BIG-IP local traffic management system, performs data compression. The default value is disabled.

`method-prefer`
Specifies the type of compression that the system prefers. The default value is gzip.

`min-size`
Specifies the minimum length in bytes of a server response that is acceptable for compression. The length in bytes applies to content length only, not headers. The default value is 1024.

`partition`
Displays the administrative partition within which the profile resides.

`selective`
Enables or disables selective compression mode. Note that the data compression feature compresses HTTP server responses, and not client requests. The default value is disabled.

`to-folder`
http-compression profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

`uri-exclude`
Disables compression on a specified list of HTTP Request-URI responses. Use a regular expression to specify a list of URIs you do not want to compress. The default value is none.

`uri-include`
Enables compression on a specified list of HTTP Request-URI responses. Use a regular expression to specify a list of URIs you want to compress. The default value is { .* }.

`vary-header`
Enables or disables the insertion of a Vary header into cacheable server responses. The default value is enabled.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2016-12-16 ltm profile http-compression(1)

ltm profile http

NAME
http - Configures an HTTP profile.

MODULE
ltm profile

SYNTAX
Configure the http component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY
create http [name]
modify http [name]
options:
accept-xff [disabled | enabled]
app-service [[string] | none]
basic-auth-realm [["string"] | none]
defaults-from [[name] | none]
description [string]
encrypt-cookie-secret [none | [passphrase]]
encrypt-cookies
[add | delete | replace-all-with] {
[cookie] ...
}
encrypt-cookies none
enforcement {
options:
rfc-compliance [disabled | enabled]
excess-client-headers [disabled | enabled]
excess-server-headers [disabled | enabled]
max-header-size [integer]
max-header-count [integer]

```

max-requests [integer]
oversize-client-headers [disabled | enabled]
oversize-server-headers [disabled | enabled]
pipeline [allow | pass-through | reject]
truncated-redirects [disabled | enabled]
unknown-method [allow | pass-through | reject]
known-methods
  [add | delete | replace-all-with] {
    [HTTP method] ...
  }
}
fallback-host [ [hostname] | none]
fallback-status-codes
  [add | delete | replace-all-with] {
[fallback status code]...
  }
fallback-status-codes none
header-erase [none | [string] ]
header-insert [none | [string] ]
insert-xforwarded-for [disabled | enabled]
lws-separator [none | string ]
lws-width [integer]
oneconnect-transformations [disabled | enabled]
oneconnect-status-reuse ["string"]
proxy-type [reverse | explicit | transparent]
redirect-rewrite [all | matching | nodes | none]
request-chunking [rechunk | sustain ]
response-chunking [rechunk | sustain | unchunk]
response-headers-permitted
  [add | delete | replace-all-with] {
[response header] ...
  }
response-headers-permitted none
server-agent-name [string]
explicit-proxy {
  options:
enabled [no | yes]
dns-resolver [dns-resolver]
ipv6 [no | yes]
tunnel-name [tunnel]
route-domain [route-domain]
default-connect-handling [deny | allow]
tunnel-on-any-request [no | yes]
connect-error-message ["string"]
dns-error-message ["string"]
bad-request-message ["string"]
bad-response-message ["string"]
}
sflow {
  options:
poll-interval [integer]
poll-interval-global [no | yes]
sampling-rate [integer]
sampling-rate-global [no | yes]
}
via-host-name [string]
via-request [append | preserve | remove]
via-response [append | preserve | remove]
xff-alternative-names
  [add | delete | replace-all-with] {
[xff alternative name] ...
  }
}
hsts {
options:
mode [enabled | disabled]
maximum-age [integer]
include-subdomains [enabled | disabled]
preload [enabled | disabled]
}
}

edit http [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties

mv http [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
options:
to-folder

reset-stats http
reset-stats http [ [ [name] | [glob] | [regex] ] ... ]

DISPLAY
list http
list http [ [ [name] | [glob] | [regex] ] ... ]
show running-config http
show running-config http [ [ [name] | [glob] | [regex] ] ... ]

```

options:
all-properties
non-default-properties
one-line
partition

show http
show http [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete http [name]

DESCRIPTION

You can use the http component to create, modify, display, or delete an HTTP profile.

The BIG-IP(r) system installation includes the following default HTTP-type profiles:

http

The default HTTP profile contains values for properties related to managing HTTP traffic.

You can create a new HTTP-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

```
create http my_http_profile defaults-from http
```

Creates a custom HTTP profile named my_http_profile that inherits its settings from the system default HTTP profile.

```
mv http /Common/my_http_profile to-folder /Common/my_folder
```

Moves a custom HTTP profile named my_http_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

Please refer to the mv manual page for examples on how to use the mv command.

OPTIONS

accept-xff

Enables or disables trusting the client IP address, and statistics from the client IP address, based on the request's XFF (X-forwarded-for) headers, if they exist.

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

basic-auth-realm

Specifies a quoted string for the basic authentication realm. The system sends this string to a client whenever authorization fails. The default value is none.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is http.

description

User defined description.

encrypt-cookie-secret

Specifies a passphrase for the cookie encryption. The default value is none.

encrypt-cookies

Specifies to encrypt specific cookies that the BIG-IP system sends to a client system. The default value is none.

enforcement

Specifies protocol enforcement options for the HTTP profile:

rfc-compliance

Specifies the behavior when non-rfc compliant traffic is seen. The default is disabled which ignores rfc non-compliance.

excess-client-headers

Specifies the pass-through behavior when max-header-count is exceeded by the client. The default is disabled which rejects the connection.

excess-server-headers

Specifies the pass-through behavior when max-header-count is exceeded by the server. The default is disabled which rejects the connection.

unknown-method

Specifies the behavior when an unknown method is seen. The default is allow which allows all methods, (known or unknown).

known-methods

Specifies the HTTP methods known by the HTTP filter. Combine with the unknown-method field to control behavior when unusual methods are parsed.

max-header-size

Specifies the maximum header size. The default value is 32768.

max-header-count

Specifies the maximum number of headers in HTTP request or response that will be handled. If client or server sends request or response with the number of headers greater than specified, the connection will be dropped. The default value is 64.

max-requests

Specifies the number of requests that the system accepts on a per-connection basis. The default value is 0 (zero), which means the system does not limit the number of requests per connection.

oversize-client-headers

Specifies the pass-through behavior when max-header-size is exceeded by the client. The default is disabled which rejects the connection.

oversize-server-headers

Specifies the pass-through behavior when max-header-size is exceeded by the server. The default is disabled which rejects the connection.

pipeline

Enables or disables HTTP/1.1 pipelining. If pass-through is chosen, then the HTTP filter will switch to pass through mode (and be disabled) if pipelined data is seen. The default value is allow, which means that clients can make requests even when prior requests have not received a response. In order for this to succeed, however, destination servers must include support for pipelining.

to-folder

http profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

truncated-redirects

Specifies the pass-through behavior when a redirect lacking the trailing carriage-return and line feed pair at the end of the headers is parsed. The default is disabled, which will silently drop the invalid HTTP.

unknown-method

Specifies the behavior (allow, reject, or pass-through) when an unknown HTTP method is parsed. The default is to allow unknown methods.

fallback-host

Specifies an HTTP fallback host. The default value is none.

With HTTP redirection, you can redirect HTTP traffic to another protocol identifier, host name, port number, or URI path. For example, if all members of a targeted pool are unavailable (that is, the members are disabled, marked as down, or have exceeded their connection limit), the system can redirect the HTTP request to the fallback host, with the HTTP reply Status Code 302 Found.

fallback-status-codes

Specifies one or more three-digit status codes that can be returned by an HTTP server. The default value is none.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

header-erase

Specifies the header string that you want to erase from an HTTP request. The default value is none.

header-insert

Specifies a quoted header string that you want to insert into an HTTP request. The default value is none.

The HTTP header being inserted can include a client IP address. Including a client IP address in an HTTP header is useful when a connection goes through a secure network address translation (SNAT) and you need to preserve the original client IP address. When you assign the configured HTTP profile to a virtual server, the system then inserts the header specified by the profile into any HTTP request that the system sends to a pool or pool member.

insert-xforwarded-for

Enables or disables insertion of an X-Forwarded-For header. The default value is disabled.

When using connection pooling, which allows clients to make use of other client requests' server connections, you can insert the X-Forwarded-For header and specify a client IP address.

lws-separator

Specifies the linear white space separator that the system uses between HTTP headers when a header exceeds the maximum width specified in the lws-width option. The valid value should be none, or, any combination of cr(carriage return), lf(line feed), or sp(space). The default value is none.

lws-width

Specifies the maximum number of columns that a header that is inserted into an HTTP request can have. The default value is 80.

name Specifies a unique name for the component. This option is required for the commands create, delete, and

modify.

oneconnect-transformations

Specifies whether the system performs HTTP header transformations for the purpose of keeping server-side connections open. The default value is enabled. This feature requires configuration of a OneConnect(tm) profile.

oneconnect-status-reuse

Specifies the 2xx and 4xx HTTP status codes that permit a server-side connection to be reused by OneConnect. The default value is "200 206". This feature requires configuration of a OneConnect(tm) profile.

partition

Displays the partition within which the component resides.

redirect-rewrite

Specifies which of the application HTTP redirects the system rewrites to HTTPS. The options are:

all Specifies to rewrite all application redirects to HTTPS.

matching

Specifies to rewrite to HTTPS only application redirects that match the original URI exactly.

nodes

If the URI contains a node IP address, instead of a host name, specifies that the system rewrites the node IP address to the virtual server IP address.

none Specifies that the system does not rewrite to HTTPS any application HTTP redirects. This is the default value.

Use this feature when an application is generating HTTP redirects that send the client to HTTP (a non-secure channel) when you want the client to continue accessing the application using HTTPS (a secure channel). This is a common occurrence when using client SSL processing on a BIG-IP system.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

request-chunking

Specifies how to handle chunked and unchunked requests. The default value is sustain. The options are described under response-chunking.

response-chunking

Specifies how to handle chunked and unchunked responses. The default value is sustain. The options are:

unchunk

If the response is chunked, this option unchunks the response, processes the HTTP content, and passes the response on as unchunked. The Keep-Alive value for the Connection header is not supported, and therefore the system sets the value of the header to close.

If the response is unchunked, the LTM system processes the HTTP content and passes the response on untouched.

rechunk

If the request or response is chunked, the system unchunks the request or response, processes the HTTP content, re-adds the chunk trailer headers, and then passes on the request or response as chunked. Any chunk extensions are lost.

If the request or response is unchunked, the system adds transfer encoding and chunking headers on egress.

sustain

Preserve request or response chunking unless there is a command to modify the body. If the request or response is chunked: unchunk the HTTP content, process the data, re-add chunking headers on egress. Chunk extensions will be lost. When the response is chunked, it can be rechunked on egress to the client.

response-headers-permitted

Specifies headers that the BIG-IP system allows in an HTTP response. The default value is none.

explicit-proxy

Specifies explicit settings for the HTTP profile:

enabled

Specifies whether the explicit proxy service is enabled or disabled. The default is no.

dns-resolver

Specifies the dns-resolver object that will be used to resolve hostnames in proxy requests. The default is dns-resolver.

ipv6 Specifies the relative order of IPv4 and IPv6 DNS resolutions for URIs. The default is no, which will try a IPv4 lookup before a IPv6.

tunnel-name

Specifies the tunnel that will be used for outbound proxy requests. This enables other virtual servers to receive connections initiated by the proxy service. The default is http-tunnel.

`route-domain`
Specifies the route-domain that will be used for outbound proxy requests. The default is 0.

`default-connect-handling`
Specifies the behavior of the proxy service for CONNECT requests. If set to deny, CONNECT requests will only be honored if there is another virtual server listening for the requested outbound connection. If set to allow outbound connections will be made regardless of other virtual servers. The default is deny.

`tunnel-on-any-request`
Specifies that the tunnel will be used for non-CONNECT requests. If set to yes, virtual servers listening on a tunnel will be able to receive any requests and default-connect-handling option effect will be extended to all outbound proxy requests. The default is no.

`host-names`
Specifies the which host names are to be treated as local. Proxy requests made for those hosts will be treated as regular HTTP requests and will be sent to the configured default pool.

`connect-error-message`
Specifies the error message that will be returned to the browser when a proxy request can't be completed because of a failure to establish the outbound connection.

`dns-error-message`
Specifies the error message that will be returned to the browser when a proxy request can't be completed because of a failure to resolve the hostname in the request.

`bad-request-message`
Specifies the error message that will be returned to the browser when a proxy request can't be completed because the request was malformed.

`bad-response-message`
Specifies the error message that will be returned to the browser when a proxy request can't be completed because the response was malformed.

`sflow`
Specifies sFlow settings for the HTTP profile:

`poll-interval`
Specifies the maximum interval in seconds between two pollings. The default value is 0. To enable this setting, you must also set the `poll-interval-global` setting to no.

`poll-interval-global`
Specifies whether the global HTTP `poll-interval` setting, which is available under `sys sflow global-settings` module, overrides the object-level `poll-interval` setting. The default value is yes.

The available values are:

no Specifies to use the object-level `poll-interval` setting.

yes Specifies to use the global HTTP `poll-interval` setting.

`sampling-rate`
Specifies the ratio of packets observed to the samples generated. For example, a sampling rate of 2000 specifies that 1 sample will be randomly generated for every 2000 packets observed. The default value is 0. To enable this setting, you must also set the `sampling-rate-global` setting to no.

`sampling-rate-global`
Specifies whether the global HTTP `sampling-rate` setting, which is available under `sys sflow global-settings` module, overrides the object-level `sampling-rate` setting. The default value is yes.

The available values are:

no Specifies to use the object-level `sampling-rate` setting.

yes Specifies to use the global HTTP `sampling-rate` setting.

`via-host-name`
Specifies the hostname that will be used in the `Via: HTTP` header. See `via-request` and `via-response` for how the `Via: header` will be handled. If either `via-request` or `via-response` are set to append, then this is required.

`via-request`
Specifies how you want to process `Via: HTTP` header in requests sent to OWS. The default setting is remove. The available values are:

append
The value from `via-host-name` is appended to the `Via: HTTP` header.

preserve
`Via: HTTP` header is preserved without changes.

remove
`Via: HTTP` header is removed from the request.

`via-response`
Specifies how you want to process `Via: HTTP` header in responses sent to clients. The default setting is

remove. The available values are the same as in via-request.

server-agent-name

Specifies the string used as the server name in traffic generated by LTM. The default value is BigIP.

alternative-xff-names

Specifies alternative XFF headers instead of the default X-forwarded-for header.

hsts Specifies HSTS settings for the HTTP profile:

mode Specifies if the HSTS settings are enabled or disabled. The default is disabled.

maximum-age

Specifies the maximum age to be sent in the HSTS header. The default is 16070400.

include-subdomains

Specifies if the includeSubdomains directive is sent in the HSTS header. The default is enabled.

preload

Specifies if the preload directive is sent in the HSTS header. The default is disabled.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2020-01-28 ltm profile http(1)

ltm profile http2

NAME

http2 - Configures a HTTP/2 protocol profile.

MODULE

ltm profile

SYNTAX

Configure the http2 component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create http2 [name]

modify http2 [name]

options:

activation-modes { [alpn | always] ... }

concurrent-streams-per-connection [integer]

connection-idle-timeout [integer]

defaults-from [[name] | none]

description [string]

frame-size [integer]

insert-header [disabled | enabled]

insert-header-name ["string"]

receive-window [integer]

write-size [integer]

header-table-size [integer]

enforce-tls-requirements [disabled | enabled]

DISPLAY

list http2

list http2 [[[name] | [glob] | [regex]] ...]

show running-config http2

show running-config http2 [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show http2

show http2 [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete http2 [name]

DESCRIPTION

You can use the http2 component to create, modify, display, or delete a HTTP/2 profile.

The BIG-IP(r) system installation includes the following default HTTP/2-type profiles:

http2

The default HTTP/2 profile contains values for properties related to managing HTTP/2 traffic.

You can create a new HTTP/2-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

```
create http2 my_http2_profile defaults-from http2
```

OPTIONS

activation-modes

Specifies what will cause a connection to be treated as a HTTP/2 connection. The value alpn specifies that the TLS application-layer-protocol-negotiation will be used to determine whether HTTP/2 should be activated. Clients that use TLS, but only support HTTP will work as-if HTTP/2 is not present. The value always specifies that all connections are assumed to be HTTP/2 connections. The default value is { alpn }.

concurrent-streams-per-connection

Specifies how many concurrent requests are allowed to be outstanding on a single HTTP/2 connection.

connection-idle-timeout

Specifies how many seconds a HTTP/2 connection is left open idly before it is shutdown.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is http2.

description

User defined description.

frame-size

Specifies the size of the data frames, in bytes, that HTTP/2 will send to the client. Larger frame sizes will improve network utilization, but may affect concurrency. The default value is 2048.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

insert-header

Specifies whether an HTTP header that indicates the use of HTTP/2 should be inserted in the request going to the back-end server. The default value is disabled.

insert-header-name

Specifies the name of the HTTP header controlled by insert-header. The default value is "X-HTTP2".

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

receive-window

Specifies the receive window, in KB. The receive window is a mechanism used by HTTP/2 to perform flow control. The receive window allows HTTP/2 to stall individual upload streams when needed. This mechanism is available only for HTTP/2 version 3. The default value is 32.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

write-size

Specifies the total size of combined data frames, in bytes, HTTP/2 will send in a single write. This controls the size of the TLS records when HTTP/2 is used over SSL. A large write size will cause HTTP/2 to buffer more data, but will improve network utilization. The default value is 16384.

header-table-size

Specifies the size of the header table, in bytes. The HTTP/2 protocol compresses http headers to save bandwidth. A larger table will allow better compression, at the cost of more memory usage. The default value is 4096 bytes.

enforce-tls-requirements

Specifies whether the TLS connection requirements, as specified in the HTTP/2 protocol specification, will be enforced. The default value is enabled.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, regex, reset-stats, show, tmsk

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2018-05-01 ltm profile http2(1)

ltm profile http3

NAME

http3 - Configures a HTTP/3 protocol profile.

MODULE

ltm profile

SYNTAX

Configure the http3 component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create http3 [name]

modify http3 [name]

options:

defaults-from [[name] | none]

description [string]

header-table-size [integer]

DISPLAY

list http3

list http3 [[[name] | [glob] | [regex]] ...]

show running-config http3

show running-config http3 [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show http3

show http3 [[[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete http3 [name]

DESCRIPTION

You can use the http3 component to create, modify, display, or delete a HTTP/3 profile.

The BIG-IP(r) system installation includes the following default HTTP/3-type profiles:

http3

The default HTTP/3 profile contains values for properties related to managing HTTP/3 traffic.

You can create a new HTTP/3-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

```
create http3 my_http3_profile defaults-from http3
```

OPTIONS

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is http3.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

header-table-size

Specifies the size of the header table, in bytes. The HTTP/3 protocol compresses http headers to save bandwidth. A larger table will allow better compression, at the cost of more memory usage. The default value is 4096 bytes.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2019-08-29 ltm profile http3(1)

ltm profile httprouter

NAME

httprouter - Configures a HTTP Router profile.

MODULE

httprouter profile

SYNTAX

Configure the httprouter component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create httprouter [name]

modify httprouter [name]

options:

app-service [(string) | none]

defaults-from [(name) | none]

description [string]

edit httprouter [[(name) | (glob) | (regex)] ...]

options:

all-properties

non-default-properties

reset-stats httprouter

reset-stats httprouter [[(name) | (glob) | (regex)] ...]

DISPLAY

list router

list router [[(name) | (glob) | (regex)] ...]

show running-config httprouter

show running-config httprouter [[(name) | (glob) | (regex)] ...]

options:

all-properties

non-default-properties

one-line

partition

show httprouter

show httprouter [[(name) | (glob) | (regex)] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete httprouter [name]

DESCRIPTION

You can use the httprouter component to manage a HTTP router profile.

EXAMPLES

```
create httprouter my_router_profile defaults-from httprouter
```

Creates a HTTP router profile named my_router_profile using the system defaults.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is router.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016, 2018. All rights reserved.

BIG-IP 2018-10-20 ltm profile httprouter(1)

ltm profile icap

NAME

icap - Configures an Internet Content Adaptation Protocol (ICAP) profile.

MODULE

ltm profile

SYNTAX

Configure the icap component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create icap [name]
modify icap [name]

options:

defaults-from [[name] | none]
description [string]
header-from [string]
host [string]
preview-length [integer]
referer [string]
uri [string]
user-agent [string]

edit icap [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

mv icap [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats icap

reset-stats icap [[[name] | [glob] | [regex]] ...]

DISPLAY

list icap

list icap [[[name] | [glob] | [regex]] ...]

show running-config icap

show running-config icap [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

show icap

show icap [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE
delete icap [name]

DESCRIPTION
You can use the icap component to manage an Internet Content Adaptation Protocol profile.

EXAMPLES
create icap my_icap defaults-from icap

Creates an internet content adaptation protocol profile named my_icap using the system defaults.

```
create icap my_icap { uri icap://mycompany.com/ad_insertion/ }
```

Creates an internet content adaptation protocol profile named my_icap that uses icap://mycompany.com/ad_insertion/ as the ICAP URI.

```
mv icap /Common/my_icap_profile to-folder /Common/my_folder
```

Moves a custom icap profile named my_icap_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is icap.

description
User defined description.

header-from
Specifies the header-from attribute to use in the ICAP header. Please refer to RFC 3507 section 4.3.2.

host Specifies the host attribute to use in the ICAP header. Please refer to RFC 3507 section 4.3.2i.

preview-length
Specifies the ICAP data preview size. Please refer to RFC 3507 section 4.5.

referer
Specifies the referer attribute to use in the ICAP header. Please refer to RFC 3507 section 4.3.2.

to-folder
icap profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

uri Specifies the ICAP URI to use in the ICAP header. Please refer to RFC 3507 section 4.2. Macro expansion has been implemented for all attributes values in the ICAP header. If an ICAP header attribute value contains \${SERVER_IP}, the macro will be replaced with the IP address of the ICAP server selected from the internal virtual server's pool. If an ICAP header attribute contains \${SERVER_PORT}, the macro will be replaced with the port of the ICAP server selected from the internal virtual server's pool. For example, the URI attribute in an ICAP profile could be set to icap://\${SERVER_IP}:\${SERVER_PORT}/videoOptimization.

user-agent
Specifies the user-agent attribute to use in the ICAP header. Please refer to RFC 3507 section 4.3.2.

SEE ALSO
create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh, ltm profile response-adapt

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012. All rights reserved.

BIG-IP 2013-12-31 ltm profile icap(1)

Itm profile iiop

NAME
iiop - Configures an Internet Inter-Orb Protocol (IIOP) profile.

MODULE
ltm profile

SYNTAX

Configure the iiop component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create iiop [name]
modify iiop [name]
options:
  abort-on-timeout [disabled | enabled]
  app-service [[string] | none]
  defaults-from [name]
  description [string]
  persist-object-key [disabled | enabled]
  persist-request-id [disabled | enabled]
  timeout [integer]
```

```
edit iiop [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

```
reset-stats iiop
reset-stats iiop [name]
```

DISPLAY

```
list iiop
list iiop [ [name] | [glob] | [regex] ] ... ]
show running-config iiop
show running-config iiop [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show iiop
show iiop [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global
```

DELETE

```
delete iiop [name]
```

DESCRIPTION

You can use the iiop component to manage IIOp network traffic. The system parses the incoming TCP stream, disaggregates it into IIOp messages, and performs load balancing and persistence based on the parameters you set.

EXAMPLES

```
create iiop my_iiop_profile defaults-from iiop
```

Creates an IIOp profile named my_iiop_profile that inherits its settings from the system default IIOp profile named iiop.

```
list iiop all-properties
```

Displays all properties for all IIOp profiles.

OPTIONS

abort-on-timeout
Specifies whether the system aborts the connection if there is no response received within the time specified in the timeout option. The default value is disabled.

app-service
Specifies the name of the application service to which the profile belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is iiop.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which this profile resides.

persist-object-key
Specifies whether to persist connections based on the object key in the IIOp request. The default value is disabled.

`persist-request-id`

Specifies whether to persist connections based on the request ID in the IOP request. The default value is enabled.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`timeout`

Specifies the request timeout. The system uses this value when the `abort-on-timeout` option is enabled. The default value is 30 seconds.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012. All rights reserved.

BIG-IP 2012-10-19 ltm profile iioip(1)

Itm profile ilx

NAME

`ilx` - Configures an ILX profile.

MODULE

ltm profile

SYNTAX

Configure the `ilx` component within the `ltm profile` module using the syntax shown in the following sections.

CREATE/MODIFY

`create ilx [name]`

`modify ilx [name]`

options:

`app-service [[string] | none]`

`defaults-from [[name] | none]`

`description [string]`

`plugin [string]`

`edit ilx [[name] | [glob] | [regex]] ...]`

options:

`all-properties`

`non-default-properties`

DISPLAY

`list ilx`

`list ilx [[name] | [glob] | [regex]] ...]`

`show running-config ilx`

`show running-config ilx [[name] | [glob] | [regex]] ...]`

options:

`all-properties`

`non-default-properties`

`one-line`

`partition`

DELETE

`delete ilx [name]`

DESCRIPTION

You can use the `ilx` component to create, modify, display, or delete an ILX profile. An ILX profile is used to associate an ILX Plugin (see `ilx plugin`) with a virtual server. When a client connects to a virtual server the data from the client and server will be routed to the ILX plugin allowing the plugin to inspect and manage traffic. An ILX profile is compatible with the following virtual server profiles: `clientssl`, `http`, `http-compression`, `serverssl`, `stream`, `tcp` and `web-acceleration`. A plugin may be associated with at most one ILX profile. An ILX profile may be associated with multiple virtual servers.

The BIG-IP(r) system installation includes the following default ILX profiles:

`ilx`

EXAMPLES

`create ilx my_ilx_profile defaults-from ilx`

Creates a custom ILX profile named `my_ilx_profile` that inherits its settings from the system default ILX profile.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is `ilx`.

`description`

User defined description.

`mpi-channel`

Displays the communication channel used by the TMM and Node.js process.

`partition`

Displays the administrative partition within which the profile resides.

`plugin`

The name of an ILX Plugin, see also `ilx plugin`.

STATS

`Connections`

The number of current active connections

`Connections Max Active`

The maximum number of active connections at one time

`Connection Aborts`

The number of aborted connections (due to plugin error)

`Connection Rejects`

The number of aborted connections (due to TMM error)

`Backlogged Messages`

Messages queued to be sent from TMM to plugin

`Backlogged Bytes`

Bytes queued to be sent from TMM to plugin

`Received Message Count`

Messages received by TMM from plugin

`Transmitted Message Count`

Messages sent from TMM to plugin

`Received Payload Bytes`

Payload bytes received by TMM from plugin

`Transmitted Payload Bytes`

Payload bytes sent from TMM to plugin

SEE ALSO

`create`, `delete`, `edit`, `glob`, `ilx plugin`, `ilx workspace`, `list`, `itm virtual`, `modify`, `regex`, `show`, `tms`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2017-05-24 itm profile ilx(1)

Itm profile imap

NAME

`imap` - Configures an IMAP profile.

MODULE

`itm profile`

SYNTAX

Configure the `imap` component within the `itm profile` module using the syntax shown in the following sections.

CREATE/MODIFY
create imap [name]
modify imap [name]
options:
 app-service [[string] | none]
 defaults-from [[name] | none]
 description [string]
 activation-mode [none | allow | require]

edit imap [[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties

DISPLAY
list imap
list imap [[name] | [glob] | [regex]] ...]
show running-config imap
show running-config imap [[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties

DELETE
delete imap [name]

DESCRIPTION
You can use the imap component to create, modify, display, or delete an IMAP profile with which you can manage IMAP traffic.

EXAMPLES
create imap my_imap_profile defaults-from imap

Creates a custom IMAP profile named my_imap_profile that inherits its settings from the system default IMAP profile.

list imap

Displays the properties of all IMAP profiles.

OPTIONS
app-service
Specifies the name of the application service to which the profile belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is imap.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

activation-mode
Sets the activation-mode STARTTLS. The options are NONE, ALLOW, or REQUIRE. The default value is REQUIRE.

SEE ALSO
create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, sys provision, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2017-08-23 ltm profile imap(1)

Itm profile ipother

NAME

ipother - Configures a generic IP profile for non-TCP and non-UDP traffic.

MODULE

itm profile

SYNTAX

Configure the ipother component within the Itm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create ipother [name]

modify ipother [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

idle-timeout [immediate | indefinite | integer]

edit ipother [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats ipother

reset-stats ipother [[[name] | [glob] | [regex]] ...]

DISPLAY

list ipother

list ipother [[[name] | [glob] | [regex]] ...]

show running-config ipother

show running-config ipother

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show ipother

show ipother [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete ipother [name]

DESCRIPTION

You can use the ipother component to manage non-TCP and non-UDP network traffic. If you want to manage TCP or UDP traffic, then use the appropriate TCP or UDP LTM profiles.

EXAMPLES

```
create ipother my_ipother_profile defaults-from ipother
```

This creates a custom IP-OTHER profile that is named my_ipother_profile which inherits its settings from the system default IP-OTHER profile.

```
list ipother all-properties
```

Displays all properties for all IP-OTHER profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is ipother.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

idle-timeout

Specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 60 seconds.

name Specifies a unique name for the component. This option is required for the commands create, delete, and

modify.

partition

Displays the administrative partition within which the profile resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, ltm profile, ltm virtual, modify, show, regex, reset-stats, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 ltm profile ipother(1)

Itm profile ipsecalg

NAME

ipsecalg - Configures a IPsecALG profile.

MODULE

ltm profile

SYNTAX

Configure the ipsecalg component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create ipsecalg [name]

modify ipsecalg [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

idle-timeout [integer]

pending-ike-connection-limit [integer]

initial-connection-timeout [integer]

log-publisher [log publisher name | none]

log-profile [log profile name | none]

mv ipsecalg [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

edit ipsecalg [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ipsecalg

list ipsecalg [[[name] | [glob] | [regex]] ...]

show running-config ipsecalg

show running-config ipsecalg [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete ipsecalg [name]

DESCRIPTION

Use this command to create, modify, display, or delete an IPsecALG profile with which you can manage IPsecALG traffic.

EXAMPLES

create ipsecalg my_ipsecalg_profile defaults-from ipsecalg

Creates a custom IPsecALG profile named my_ipsecalg_profile that inherits its settings from the system default IPsecALG profile.

list ipsecalg

Displays the properties of all IPsecALG profiles.

mv ipsecalg /Common/my_ipsecalg_profile to-folder /Common/my_folder

Moves a custom IPsecALG profile named my_ipsecalg_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is ipsecalg.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

to-folder

ipsecalg profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

idle-timeout

Specifies an idle timeout in seconds. This setting specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 3600 seconds.

pending-ike-connection-limit

Specifies the the maximum number of unacknowledged IKE connections a client can have before being denied further requests to prevent a single client from flooding all the connections trying to establish the connections. The default value is 5 connections per client.

initial-connection-timeout

Specifies an initial connection timeout in seconds. This is the max number of seconds to wait for a response from the server for the IKE/IPsec request. The default value is 3 seconds.

log-publisher

Specify the name of the log publisher which logs translation events. See help sys log-config for more details on the logging sub-system. Use the "sys log-config publisher" component to set up a log publisher.

log-profile

Specify the name of the ALG log profile which controls the logging of ALG events. See help ltm alg-log-profile for more details on the logging profile sub-system. Use the "ltm alg-log-profile profile" component to set up an ALG log profile.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014, 2016. All rights reserved.

BIG-IP 2017-01-20 ltm profile ipsecalg(1)

ltm profile mapt

NAME

map-t - Configures a MAP-T tunnel profile.

MODULE

net tunnels

SYNTAX

Configure the map-t component within the ltm profile module using the syntax in the following sections.

CREATE/MODIFY

create map-t [name]

modify map-t [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

ip6-prefix [ipv6 address/netmask]

ip4-prefix [ipv4 address/netmask]

ea-bits-length [integer]

port-offset [integer]

br-prefix [ipv6 address/netmask]

edit map-t [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list map-t

list map-t [[[name] | [glob] | [regex]] ...]

show running-config map-t

show running-config map-t [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete map-t [all | [name]]

DESCRIPTION

You can use the map-t component to create a MAP-T profile that you associate with an LTM Virtual. A map-t profile allows you to define NAT parameters which will allow legacy IPV4 servers to connect to an IPV6 network.

EXAMPLES

```
create map-t my_map
```

Creates a MAP-T profile called my_map.

```
list map-t all-properties
```

Displays all the properties of all the MAP-T profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is map.t.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ip6-prefix

Specifies the Rule IPv6 Prefix and netmask using CIDR notation, such as 2014::/48. The default prefix length is 48.

ip4-prefix

Specifies the Rule IPv4 Prefix and netmask using CIDR notation, such as 192.0.0.0/8. The default prefix length is 8.

ea-bits-length

Specifies the Rule EA (Embedded Address) Length of the MAP-T domain. The default is 40 (IPv4 32 bits + PSID 8 bits).

port-offset

Specifies the port offset bits length of the MAP-T domain. The default is 6.

br-prefix

Specifies the BR (Border Relay) IPv6 prefix and netmask using CIDR notation, such as 2023::/96. The default prefix length is 96.

psid Specify PSID value for MAPT. The default is zero.

draft-mode

Allow BIGIP to work with non rfc compliant CE devices which map BR prefix at a wrong offset. The default is to disable draft mode.

br-pfx-ignore

Allow BIGIP to work with non rfc compliant CE devices which ignore BR prefix lengths. The default is to disable.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2014, 2016. All rights reserved.

BIG-IP 2017-01-20 ltm profile mapt(1)

ltm profile mblb

NAME

mblb - Configures an MBLB profile (experimental).

MODULE

ltm profile

SYNTAX

Configure the mblb component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create mblb [name]

modify mblb [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

isolate-abort [disabled | enabled]

isolate-expire [disabled | enabled]

isolate-server [disabled | enabled]

isolate-client [disabled | enabled]

egress-high [# of messages]

egress-low [# of messages]

ingress-high [# of messages]

ingress-low [# of messages]

min-conn [# of connections]

tag-ttl [# of seconds]

shutdown-timeout [# of seconds]

edit mblb [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list mblb

list mblb [[name] | [glob] | [regex]] ...]

show running-config mblb

show running-config mblb [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DELETE

delete mblb [name]

DESCRIPTION

Use this command to create, modify, display, or delete an MBLB profile with which you can customize MBLB behavior.

EXAMPLES

```
create mblb my_mblb_profile defaults-from mblb
```

Creates a custom MBLB profile named my_mblb_profile that inherits its settings from the system default MBLB profile.

```
list mblb
```

Displays the properties of all MBLB profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is mblb.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

isolate-abort

Specify whether to isolate abort event propagation.

isolate-expire

Specify whether to isolate expiration event propagation.

isolate-server

Specify whether to isolate serverside shutdown event propagation. This also dominates serverside abort/expiration event propagation.

isolate-client

Specify whether to isolate clientside shutdown event propagation. This also dominates clientside abort/expiration event propagation.

egress-high

Specify the high water mark for egress message queue. The default value is 50.

egress-low

Specify the low water mark for egress message queue. The default value is 5.

ingress-high

Specify the high water mark for ingress message queue. The default value is 50.

ingress-low

Specify the low water mark for ingress message queue. The default value is 5.

min-conn

Specify the minimum number of serverside connections. The default value is 0.

tag-ttl

Specify the TTL (time to live) for message TAG. The default value is 60.

shutdown-timeout

Delays sending FIN when BIGIP receives the first FIN packet from either the client or the server. Value of 0 means send FIN immediately otherwise the minimum of tcp idle timeout and shutdown timeout is used. The default value is 5 seconds

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, sys provision, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014. All rights reserved.

ltm profile mqtt

NAME

mqtt - Configures an MQTT profile.

MODULE

ltm profile

SYNTAX

Configure the mqtt component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create mqtt [name]

modify mqtt [name]

options:

app-service [[string] | none]

edit mqtt [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv mqtt [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats mqtt

reset-stats mqtt [[[name] | [glob] | [regex]] ...]

DISPLAY

list mqtt

list mqtt [[[name] | [glob] | [regex]] ...]

show running-config mqtt

show running-config mqtt [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show mqtt

show mqtt [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete mqtt [name]

DESCRIPTION

You can use the mqtt component to create, modify, display, or delete an MQTT profile.

The BIG-IP(r) system installation includes the following default MQTT-type profiles:

mqtt

The default MQTT profile contains values for properties related to managing MQTT traffic.

You can create a new MQTT-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

```
create mqtt my_mqtt_profile defaults-from mqtt
```

Creates a custom MQTT profile named my_mqtt_profile that inherits its settings from the system default MQTT profile.

```
mv mqtt /Common/my_mqtt_profile to-folder /Common/my_folder
```

Moves a custom MQTT profile named my_mqtt_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

Please refer to the mv manual page for examples on how to use the mv command.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-09-05 ltm profile mqtt(1)

ltm profile mssql

NAME

mssql - Configures a profile to manage mssql(tds) database traffic.

MODULE

ltm profile

SYNTAX

Configure the mssql component within the ltm profile module using the syntax in the following sections.

CREATE/MODIFY

create mssql [name]

modify mssql [name]

options:

app-service [[string] | none]

defaults-from [name]

description [[string] | none]

read-pool [string]

read-write-split-by-user [disabled | enabled]

read-write-split-by-command [disabled | enabled]

user-can-write-by-default [true | false]

user-list [add | delete | none | replace-all-with] {
[user names...]

}

write-persist-timer [number]

write-pool [string]

edit mssql [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats mssql

reset-stats mssql [[[name] | [glob] | [regex]] ...]

DISPLAY

list mssql

list mssql [[[name] | [glob] | [regex]] ...]

show running-config mssql

show running-config mssql [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show mssql

show mssql [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete mssql [name]

DESCRIPTION

You can use the mssql component to configure a profile to manage mssql(tds) database traffic.

EXAMPLES

create mssql my_mssql_profile defaults-from mssql

Creates a mssql profile named my_mssql_profile that inherits its settings from the system default mssql profile.

list mssql

Displays the properties of all mssql profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is mssql.

partition

Displays the administrative partition within which the profile resides.

read-pool

Specifies the pool of MS SQL database servers to which the system sends ready-only requests.

read-write-split-by-command

When enabled, the system decides which pool to send the client requests the by the content in the message. It can only be enabled when read-write-split-by-user is disabled.

read-write-split-by-user

When enabled, the system decides which pool to send the client requests the by user name. It can only be enabled when read-write-split-by-command is disabled.

user-can-write-by-default

Specifies whether users have write access by default. When set to true, all users have write access, except those added to the users list.

user-list

Specifies the users who have read-only access to the MS SQL if user-can-write-by-default is true, or the users who have write access to the MS SQL database if user-can-write-by-default is false.

write_persist_timer

Specify how many minimum time in milliseconds the connection will be persisted to write-pool after connection switch to write pool.

write-pool

Specifies the pool of MS SQL database servers to which the system sends requests that are not read-only.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2014-09-25 ltm profile mssql(1)

ltm profile netflow

NAME

netflow - Configures a NETFLOW profile.

MODULE

ltm profile

SYNTAX

Configure the netflow component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create netflow [name]

modify netflow [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

netflow-version [enumerated]

edit netflow [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

mv netflow [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats netflow

reset-stats netflow [[[name] | [glob] | [regex]] ...]

DISPLAY

list netflow

list netflow [[[name] | [glob] | [regex]] ...]

show running-config netflow

show running-config netflow [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

show netflow

show netflow [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE

delete netflow [name]

DESCRIPTION

You can use the netflow component to manage a NETFLOW profile.

EXAMPLES

```
create netflow my_netflow_profile defaults-from netflow
```

Creates a NETFLOW profile named my_netflow_profile using the system defaults.

```
create netflow my_netflow_profile { netflow-version version-5 }
```

Creates a NETFLOW profile named my_netflow_profile that specifies the maximum number of messages that can be held in the ingress queue is 1000.

```
mv netflow /Common/my_netflow_profile to-folder /Common/my_folder
```

Moves a custom netflow profile named my_netflow_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is netflow.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

netflow-version

Specifies the expected netflow version for a profile.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

ltm profile ntlm

NAME

ntlm - Configures a Microsoft(r) Windows(r) NT Local Area Network (LAN) manager profile.

MODULE

ltm profile

SYNTAX

Configure the ntlm component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create ntlm [name]

modify ntlm [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

insert-cookie-domain [domain]

insert-cookie-name [cookie name]

insert-cookie-passphrase [passphrase]

key-by-cookie [disabled | enabled]

key-by-cookie-name [cookie name]

key-by-domain [disabled | enabled]

key-by-ip-address [disabled | enabled]

key-by-target [disabled | enabled]

key-by-user [disabled | enabled]

key-by-workstation [disabled | enabled]

edit ntlm [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ntlm

list ntlm [[[name] | [glob] | [regex]] ...]

show running-config ntlm

show running-config ntlm [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete ntlm [name]

DESCRIPTION

You can use the ntlm component to create a Microsoft Windows NT LAN manager (NTLM) profile to manage servers on the LAN that are running Windows NT.

EXAMPLES

```
create ntlm my_ntlm_profile defaults-from ntlm
```

Creates a Microsoft Windows NT LAN manager profile named my_ntlm_profile that inherits its settings from the system default NTLM profile named ntlm.

```
list ntlm all-properties
```

Displays all properties for all NTLM profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is ntlm.

description

User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`insert-cookie-domain`
Specifies an optional domain for the inserted cookie. The default is none, which causes no domain to be configured for the inserted cookie.

`insert-cookie-name`
Specifies a cookie name that the system inserts in the cookie. The default value is `NTLMconnpool`.

`insert-cookie-passphrase`
Specifies a cookie passphrase that the system inserts in the cookie. The default value is `mypassphrase`.

`key-by-cookie`
Specifies whether the system uses the existing cookie as the key. The default value is disabled.

`key-by-cookie-name`
Specifies whether the system uses the value of the `insert-cookie-name` option as the key. The default value is `mycookie`.

`key-by-domain`
Specifies whether the system uses the NTLM domain as the key. The default value is disabled.

`key-by-ip-address`
Specifies whether the system uses the client IP address as the key. The default value is disabled.

`key-by-target`
Specifies whether the system uses the NTLM target as the key. The default value is disabled.

`key-by-user`
Specifies whether the system uses the NTLM user as the key. The default value is enabled.

`key-by-workstation`
Specifies whether the system uses the NTLM workstation as the key. The default value is disabled.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`regex`
Displays the items that match the regular expression. The regular expression must be preceded by an `@` sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012. All rights reserved.

BIG-IP 2012-10-19 ltm profile ntlm(1)

ltm profile ojsp-stapling-params

NAME

`ojsp-stapling-params` - NOTICE: The use of `ltm profile ojsp-stapling-params` has been deprecated since v13.0.0. Please use `sys crypto cert-validator ojsp` instead.

MODULE

`ltm profile`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2017-01-20 ltm profile ojsp-stapling-params(1)

Itm profile ocspp

NAME

ocsp - Configures a OCSP profile.

MODULE

itm profile

SYNTAX

Configure the ocsp component within the Itm profile module using the syntax in the following sections.

CREATE/MODIFY

create ocsp [name]

modify ocsp [name]

options:

defaults-from [[name] | none]

max-age [integer]

nonce [enabled | disabled]

edit ocsp [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv ocsp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list ocsp

list ocsp [[name] | [glob] | [regex]] ...]

show running-config ocsp

show running-config ocsp

[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete ocsp [all | [name]]

options:

recursive

DESCRIPTION

You can use the ocsp component to manage a OCSP profile.

OCSP profile enables the traffic management system to respond to OCSP requests from clients. You can implement this type of profile by using the default profile, or by creating a custom profile based on the OCSP profile template and modifying its settings.

EXAMPLES

```
create ocsp my_ocsp_profile defaults-from ocsp
```

Creates a custom OCSP profile named my_ocsp_profile that inherits its settings from the system default profile ocsp.

```
list ocsp all-properties
```

Displays all properties for all OCSP profiles.

```
mv ocsp /Common/my_ocsp_profile to-folder /Common/my_folder
```

Moves a custom ocsp profile named my_ocsp_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is ocsp.

max-age

Specifies the value for HTTP Response cache control header max-age. The max-age header sent to client is the lesser of the configured value and validity of OCSP response. Default value is 604800 seconds.

nonce

Specifies whether OCSP Nonce Request Extension is supported by the OCSP profile. Default value is enabled.

partition

Displays the administrative partition within which the component resides.

SEE ALSO

create, delete, edit, glob, list, ltm profile client-ssl, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2019-2020. All rights reserved.

BIG-IP 2020-01-29 ltm profile ocsp(1)

ltm profile one-connect

NAME

one-connect - Configures a OneConnect(tm) profile.

MODULE

ltm profile

SYNTAX

Configure the one-connect component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create one-connect [name]

modify one-connect [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

idle-timeout-override [disabled | enabled]

share-pools [disabled | enabled]

max-age [integer]

max-reuse [integer]

max-size [integer]

source-mask [ip address]

limit-type [none|idle|strict]

edit one-connect [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv one-connect [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats one-connect

reset-stats one-connect [[[name] | [glob] | [regex]] ...]

DISPLAY

list one-connect

list one-connect [[[name] | [glob] | [regex]] ...]

show running-config one-connect

show running-config one-connect [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show one-connect

show one-connect [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete one-connect [name]

DESCRIPTION

You can use the one-connect component to create a OneConnect profile that optimizes connections by improving client performance and increasing server capacity.

EXAMPLES

create one-connect my_OC_profile defaults-from oneconnect

Creates a OneConnect profile named my_OC_profile that inherits its settings from the system default OneConnect profile named oneconnect.

list one-connect all-properties

Displays all properties for all OneConnect profiles.

mv one-connect /Common/my_oneconnect_profile to-folder /Common/my_folder

Moves a custom one-connect profile named my_oneconnect_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is oneconnect.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

idle-timeout-override

Specifies the number of seconds that a connection is idle before the connection flow is eligible for deletion. The default value is disabled.

share-pools

Indicates that connections may be shared not only within a virtual server, but also among similar virtual servers (e.g. those that differ only in destination address). When enabled, all virtual servers that use the same One Connect and other internal network profiles can share connections.

max-age

Specifies the maximum age, in number of seconds, of a connection in the connection reuse pool. For any connection with an age higher than this value, the system removes that connection from the reuse pool. The default value is 86400.

max-reuse

Specifies the maximum number of times that a server connection can be reused. The default value is 1000.

max-size

Specifies the maximum number of connections that the system holds in the connection reuse pool. If the pool is already full, then the server connection closes after the response is completed. The default value is 10000.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

source-mask

Specifies a source IP mask. The default value is 0.0.0.0.

The system applies the value of this option to the source address to determine its eligibility for reuse. A mask of 0.0.0.0 causes the system to share reused connections across all clients. A host mask (all 1's in binary), causes the system to share only those reused connections originating from the same client IP address.

limit-type

Connection limits with OneConnect are different from straight TCP connection limits. Three options are supported: "none" (the default), "idle", and "strict". When the limit is "none", simultaneous in-flight requests and responses over TCP connections to a pool member are counted toward the limit (this being the historical handling). There may be more TCP connections open to support new requests than there can be simultaneous in-flight requests and responses. This is particularly true when SNAT pools and narrow source address masks are used. When the limit is "idle", idle connections will be dropped as the TCP connection limit is reached. For short intervals, during the overlap of the idle connection being dropped and the new connection being established, the TCP connection limit may be exceeded. When the limit is "strict", the TCP connection limit is honored with no exceptions. This means that idle connections will prevent new TCP connections from being made until they expire, even if they could otherwise be reused. This is not a recommended configuration except in very special cases with short expiration timeouts.

to-folder

one-connect profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2014-03-31 ltm profile one-connect(1)

ltm profile pcp

NAME

pcp - Configures a PCP profile.

MODULE

ltm profile

SYNTAX

CREATE/MODIFY

create pcp [name]

modify pcp [name]

options:

announce-after-failover [enabled | disabled]

announce-multicast [integer]

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

map-filter-limit [integer]

map-limit-per-client [integer]

map-recycle-delay [integer]

max-mapping-lifetime [integer]

min-mapping-lifetime [integer]

rule [[rule_name] | none]

third-party-allowed-subnets

[add | delete | replace-all-with] {
[ip address/prefix length] ...
}

third-party-option [enabled | disabled]

edit pcp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list pcp

list pcp [[[name] | [glob] | [regex]] ...]

show running-config pcp

show running-config pcp

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show pcp

show pcp [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete pcp [name]

DESCRIPTION

You can use the pcp component to specify Port Control Protocol attributes for a profile that can be used in an LSN pool.

EXAMPLES

```
create pcp my_pcp_profile defaults-from pcp
```

Creates a custom PCP profile named my_pcp_profile that inherits its settings from the system default pcp profile.

list pcp all-properties

Displays all properties for all PCP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is pcp, a profile that is shipped in the software.

description

User defined description.

announce-after-failover

Specifies that the BIG-IP software should send an unsolicited ANNOUNCE response to all PCP clients when there is a failover. The unsolicited ANNOUNCE response goes over a link-local multi-cast address, and it contains a new EPOCH time. This signals to the PCP clients that they should renew all of their active mappings.

announce-multicast

Whenever the BIG-IP system reboots, or if there is any possibility that the system lost its PCP-mapping state, it sends an unsolicited ANNOUNCE response to all of its PCP clients. It sends the response over a link-local multi-cast address, and it contains a new EPOCH time. The PCP clients react by renewing all of their active IP mappings. To compensate for possible packet loss (since the multi-cast address is link-local), you can use this property to set the number of multi-cast re-sends. Default is 10 re-sends.

map-filter-limit

A PCP client can request a "filter" for a mapping entry, where the filter limits the number of external endpoints that can use the IP map. The filter request contains the particular IP address and port for the endpoint (or subnet of endpoints), as well as a prefix length. Enter the maximum number of filters (allowed subnets) that clients are allowed to set for each PCP mapping. Default is 1.

map-limit-per-client

Specifies the maximum number of PCP mappings per client. Default is 65535 (unlimited).

Use run util lsndb to see the currently-active set of PCP mappings on the system. See "util lsndb" for details on the LSN DB utility.

map-recycle-delay

After a IP mapping times out (that is, its lifetime expires), there is a further delay before the public-side address and port can be used by another PCP client. Use this property to set the recycle delay. Default is 60 (seconds).

Use run util lsndb to see the currently-active set of PCP mappings on the system. See "util lsndb" for details on the LSN DB utility.

max-mapping-lifetime

When a PCP client requests an IP mapping from a BIG IP system, it also requests a "lifetime" for the mapping. The mapping expires at the end of that lifetime. This property is the maximum number of seconds allowed for a mapping lifetime. Default is 86400 (seconds), or 1 day.

Use run util lsndb to see the currently-active set of PCP mappings on the system. See "util lsndb" for details on the LSN DB utility.

min-mapping-lifetime

Specifies the minimum number of seconds allowed for a mapping lifetime. Default is 600 (seconds), or 10 minutes.

Use run util lsndb to see the currently-active set of PCP mappings on the system. See "util lsndb" for details on the LSN DB utility.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex (regex) for a description of regular expression syntax.

rule Specifies the iRule that is associated with this pcp profile. An iRule can read packets and possibly filter them based on whatever programming logic you design. For example, an iRule could reject all PCP mapping requests using a specific port, or pass an ANNOUNCE request through a specific port. An iRule gives you the flexibility to filter, process, or log the PCP packets that fit this profile.

Select an iRule from the menu of existing iRules. To create a new one, use the create ltm rule command (see "ltm rule").

third-party-allowed-subnets

Specifies the PCP clients that can make MAP requests on behalf of other clients. Enter a collection of IP prefixes (IPv4 or IPv6) with their prefix lengths. If a PCP client outside of any of these subnets attempts a PCP mapping, the BIG-IP software rejects the mapping.

You can shorten any IPv6 addresses as defined in RFC 2373 (see).

This list is only used if the third-party-option is also enabled.

If the list is empty and the third-party-option is enabled, any PCP client can create mappings for third parties.

third-party-option

Allows PCP clients to make MAP requests on behalf of other clients, using the THIRD_PARTY flag in the PCP request. You can set this property to enabled or disabled. If you enable this property, we recommend using the third-party-subnets option to limit the clients that can use the THIRD_PARTY flag; it is a potential security risk. The default is disabled.

SEE ALSO

create, delete, edit, list, ltm lsn-pool, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm profile pcp(1)

ltm profile pop3

NAME

pop3 - Configures a POP3 profile.

MODULE

ltm profile

SYNTAX

Configure the pop3 component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create pop3 [name]
modify pop3 [name]
```

options:

```
app-service [[string] | none]
defaults-from [ [name] | none]
description [string]
activation-mode [ none | allow | require ]
```

```
edit pop3 [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list pop3
list pop3 [ [ [name] | [glob] | [regex] ] ... ]
show running-config pop3
show running-config pop3 [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DELETE

```
delete pop3 [name]
```

DESCRIPTION

You can use the pop3 component to create, modify, display, or delete a POP3 profile with which you can manage POP3 traffic.

EXAMPLES

```
create pop3 my_pop3_profile defaults-from pop3
```

Creates a custom POP3 profile named my_pop3_profile that inherits its settings from the system default POP3 profile.

```
list pop3
```

Displays the properties of all POP3 profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is pop3.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

activation-mode

Sets the activation-mode for STARTTLS. The options are NONE, ALLOW, or REQUIRE. The default value is REQUIRE.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, sys provision, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2017-08-22 ltm profile pop3(1)

ltm profile pptp

NAME

pptp - Configures a Point-to-Point Tunneling Protocol (PPTP) profile.

MODULE

ltm profile

SYNTAX

Configure the pptp component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create pptp [name]

modify pptp [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [[string] | none]

publisher-name [[string] | none]

include-destination-ip [disabled | enabled]

csv-format [disabled | enabled]

edit pptp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats pptp

reset-stats pptp [[[name] | [glob] | [regex]] ...]

DISPLAY

list pptp

list pptp [[[name] | [glob] | [regex]] ...]

show running-config pptp

show running-config pptp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show pptp

show pptp [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE
delete pptp [name]

DESCRIPTION
You can use the pptp component to manage a PPTP profile.

EXAMPLES
create pptp my_pptp_profile defaults-from pptp
Creates a PPTP profile named my_pptp_profile using the system defaults.
create pptp my_pptp_profile { log-server-ip disabled }
Creates a PPTP profile named my_pptp_profile with server address logging disables.
modify pptp my_pptp_profile description "This is a PPTP Profile"
Modifies the description attribute of a PPTP profile named my_pptp_profile.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is pptp.
description
User defined description.
glob
Displays the items that match the glob expression. See help glob for a description of glob expression syntax.
publisher-name
Specifies the name of the log publisher for PPTP events.
include-destination-ip
Specifies whether the log messages for call establishment/disconnect include the server's ip address. The default value is disabled. When disabled the ip address will be displayed as 0.0.0.0.
csv-format
When enabled, use CSV log format for log entries. The default value is disabled.
name
Specifies a unique name for the component. This option is required for the commands create, delete, and modify.
partition
Displays the administrative partition within which the component resides.
regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO
create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013. All rights reserved.

BIG-IP 2016-03-09 ltm profile pptp(1)

ltm profile qoe

NAME
qoe - Deprecated since v15.0.0. Configures a Quality of Experience (QoE) Monitoring profile.

MODULE
ltm profile

SYNTAX
Configure the qoe component within the ltm profile module using the syntax shown in the following sections.

```
CREATE/MODIFY
create qoe [name]
modify qoe [name]
options:
  app-service [[string] | none]
  defaults-from [name]
  description [[string] | none]
  video [true | false]

reset-stats qoe
reset-stats qoe [ [ [name] | [glob] | [regex] ] ... ]

DISPLAY
list qoe
list qoe [ [ [name] | [glob] | [regex] ] ... ]
show running-config qoe
show running-config qoe [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

show qoe
show qoe [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt

DELETE
delete qoe [name]
```

DESCRIPTION
You can use the qoe component to monitor Video Quality of Experience.

EXAMPLES
create qoe my_qoe defaults-from qoe

Creates an quality of experience profile named my_qoe.

create qoe my_qoe { video true }

video
Specifies to monitor the QoE MOS score of video streams with the format of MP4 or FLV.

SEE ALSO
create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh, ltm profile qoe

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013. All rights reserved.

BIG-IP 2018-10-31 ltm profile qoe(1)

Itm profile quic

NAME
quic - Configures a Quic profile.

MODULE
ltm profile

SYNTAX
Configure the quic component within the ltm profile module using the syntax shown in the following sections.

```
CREATE/MODIFY
create quic [name]
modify quic [name]
options:
  defaults-from [[name] | none]
  description [string]
  bidi-concurrent-streams-per-connection [integer]
  uni-concurrent-streams-per-connection [integer]

edit quic [ [ [name] | [glob] | [regex] ] ... ]
```

options:
all-properties
non-default-properties

mv quic [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]
options:
to-folder

reset-stats quic
reset-stats quic [[[name] | [glob] | [regex]] ...]

DISPLAY
list quic
list quic [[[name] | [glob] | [regex]] ...]
show running-config quic
show running-config quic
[[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

show quic
show quic [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete quic [name]

DESCRIPTION
You can use the quic component to manage Quic network traffic.

EXAMPLES
create quic my_quic_profile defaults-from quic

Creates a custom Quic profile named my_quic_profile that inherits its settings from the system default Quic profile.

list quic all-properties

Displays all properties for all Quic profiles.

mv quic /Common/my_quic_profile to-folder /Common/my_folder

Moves a custom Quic profile named my_quic_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS
app-service
Specifies the name of the application service to which the profile belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is quic.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

bidi-concurrent-streams-per-connection
Specifies how many bidirectional concurrent streams are allowed to be outstanding on a single QUIC connection.

uni-concurrent-streams-per-connection
Specifies how many unidirectional concurrent streams are allowed to be outstanding on a single QUIC connection.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the profile resides.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

to-folder

quic profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, ltm profile, ltm virtual, modify, mv, show, regex, reset-stats, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2018. All rights reserved.

BIG-IP 2019-03-04 ltm profile quic(1)

ltm profile radius

NAME

radius - Configures a RADIUS profile for network traffic load balancing.

MODULE

ltm profile

SYNTAX

Configure the radius component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create radius [name]

modify radius [name]

options:

app-service [[string] | none]

clients [add | delete | modify | replace-all-with] {
[ip address] ...
}

clients none

defaults-from [name]

description [string]

persist-avp [[string] | [integer] | none]

pem-protocol-profile-radius [[pem_protocol_profile_radius_name] | none]

subscriber-discovery [disabled | enabled]

mv radius [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

edit radius [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats radius

reset-stats radius [[[name] | [glob] | [regex]] ...]

DISPLAY

list radius

list radius [[[name] | [glob] | [regex]] ...]

show running-config radius

show running-config radius [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show radius

show radius [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete radius [name]

DESCRIPTION

You can use the radius component to manage RADIUS network traffic.

EXAMPLES

create radius my_radius_server

Creates a RADIUS profile named `my_radius_server` that inherits its settings from the system default RADIUS profile.

```
delete radius my_radius_server
```

Deletes the RADIUS profile named `my_radius_server`.

```
mv radius /Common/my_radius_profile to-folder /Common/my_folder
```

Moves a custom radius profile named `my_radius_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`clients`

Specifies host and network addresses from which clients can connect. The default value is none, which indicates that any client can connect.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is `radiusLB`.

`description`

User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`pem-protocol-profile-radius`

Specifies PEM protocol profile to be used when subscriber discovery is enabled. PEM protocol profile defines mapping of RADIUS AVPs to subscriber ID and other PEM subscriber session attributes.

`persist-avp`

Specifies the name of the RADIUS attribute on which traffic persists. Acceptable values are ASCII strings from section 5 of RFC 2865 or numeric codes (1-255). The default value is none, which indicates that persistence is disabled.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an `@` sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`subscriber-discovery`

Specifies whether to enable PEM subscriber discovery based on the content of RADIUS packets. The options are `disabled` and `enabled`. The default value is `disabled`, which indicates that it will not extract subscriber information from RADIUS packets.

`to-folder`

radius profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `mv`, `regex`, `reset-stats`, `show`, `tmssh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm profile radius(1)

ltm profile ramcache

NAME

`ramcache` - Manages the BIG-IP(r) system RAM cache.

MODULE

ltm profile

SYNTAX

Configure the ramcache component within the ltm profile module using the syntax shown in the following sections.

DISPLAY

```
show ramcache
show ramcache [ [ [name] | [glob] | [regex] ] ... ]
options:
  exact
  host [string]
  max-response [integer]
  uri [string]
```

DELETE

```
delete ramcache [name]
```

DESCRIPTION

You can use the ramcache component to delete the entries in or show information about the BIG-IP(r) system RAM cache.

EXAMPLES

```
show ramcache
```

Displays information about the entries in the BIG-IP system RAM cache.

```
delete ramcache
```

Deletes the entries in the BIG-IP system RAM cache.

OPTIONS

exact

Displays the exact number of entries in the RAM cache.

host Displays the host from which the entry was cached.

max-response

Displays the maximum number of entries in the RAM cache. The default value is 0 (zero), which is equivalent with no max-response value being specified. Without the max-response option the system will limit the number of entries to 10 per Traffic Management Microkernel (TMM).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

uri Displays the URI from which the entry was cached.

SEE ALSO

delete, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2016-10-10 ltm profile ramcache(1)

ltm profile request-adapt

NAME

request-adapt - Configures a HTTP request adaptation profile.

MODULE

ltm profile

SYNTAX

Configure the request-adapt component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create request-adapt [name]
modify request-adapt [name]
options:
  defaults-from [ [name] | none]
```

enabled [yes | no]
internal-virtual [[name] | none]
preview-size [integer]
service-down-action [ignore | reset | drop]
timeout [integer]
allow-http-10 [yes | no]

edit request-adapt [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

mv request-adapt [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats request-adapt

reset-stats request-adapt [[[name] | [glob] | [regex]] ...]

DISPLAY

list request-adapt

list request-adapt [[[name] | [glob] | [regex]] ...]

show running-config request-adapt

show running-config request-adapt [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

show request-adapt

show request-adapt [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE

delete request-adapt [name]

DESCRIPTION

You can use the request-adapt component to manage a HTTP request adaptation profile.

EXAMPLES

```
create request-adapt my_req_adapt defaults-from request-adapt
```

Creates a HTTP request adaptation profile named my_req_adapt using the system defaults.

```
create request-adapt my_req_adapt { enabled yes }
```

Creates a HTTP request adaptation profile named my_req_adapt that is enabled for adapting HTTP requests.

```
mv request-adapt /Common/my_requestadapt_profile to-folder /Common/my_folder
```

Moves a custom request-adapt profile named my_requestadapt_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is requestadapt.

enabled

Enables adaptation of HTTP requests. If set to yes, HTTP requests will be forwarded to the specified internal virtual server for adaptation. The default value is yes.

internal-virtual

Specifies the name of the internal virtual server to use for adapting the HTTP request.

preview-size

Specifies the maximum size of the preview buffer. The preview buffer is used to hold a copy of the HTTP request header and data sent to the internal virtual server in case the adaptation server reports that it does not need to adapt the HTTP request. Setting the preview-size to 0, disables buffering the request and should only be done if the adaptation server will always return with a modified HTTP request or the original HTTP request. The default value is 1024.

service-down-action

Specifies the action to take if the internal virtual server does not exist or returns an error. The default value is ignore.

The options are:

ignore

Ignore the error and send the unmodified HTTP request to a HTTP server selected from this virtual server's pool.

drop Drop the connection.

reset

Reset the connection.

timeout

Specifies a timeout in milliseconds. If the internal virtual server does not return a result within the specified time, a timeout error will occur. If preview-size is 0, the countdown will start after the entire content has been sent for adaptation. A 0 value disables the timeout. The default value is 0.

to-folder

request-adapt profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

allow-http-10

Specifies whether to forward HTTP version 1.0 requests for adaptation. By default only HTTP version 1.1 requests are forwarded. Version 1.0 is not supported. While it should work in most cases, it might be necessary to restrict adaptation on a site-specific basis. The default value is no.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh, ltm profile
response-adapt

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2017-05-24 ltm profile request-adapt(1)

ltm profile request-log

NAME

request-log - Configures a Request-Logging profile.

MODULE

ltm profile

SYNTAX

Configure the request-log component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create request-log [name]

modify request-log [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

log-request-logging-errors [disabled | enabled]

log-response-by-default [disabled | enabled]

log-response-logging-error [disabled | enabled]

proxy-close-on-error [disabled | enabled]

proxy-respond-on-logging-error [disabled | enabled]

proxy-response [string]

request-log-error-pool [[pool_name] | none]

request-log-error-protocol [TCP | UDP | none]

request-log-error-template [string]

request-log-pool [[pool_name] | none]

request-log-protocol [TCP | UDP | none]

request-log-template [string]

request-logging [disabled | enabled]

response-log-error-pool [[pool_name] | none]

response-log-error-protocol [TCP | UDP | none]

response-log-error-template [string]

response-log-pool [[pool_name] | none]

response-log-protocol [TCP | UDP | none]

response-log-template [string]

response-logging [disabled | enabled]

edit request-log [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv request-log [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

```
list request-log
list request-log [ [ [name] | [glob] | [regex] ] ... ]
show running-config request-log
show running-config request-log
[ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show request-log
show request-log [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global
```

```
DELETE
delete request-log [name]
```

DESCRIPTION

You can use the request-log component to manage request-log network traffic.

EXAMPLES

```
create request-log my_reqlog_profile defaults-from request-log
```

Creates a custom request-log profile named my_reqlog_profile that inherits its settings from the system default request-log profile.

```
list request-log all-properties
```

Displays all properties for all request-log profiles.

```
mv request-log /Common/my_requestlog_profile to-folder /Common/my_folder
```

Moves a custom request-log profile named my_requestlog_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the default values from this profile.

description

User defined description.

log-request-logging-errors

Enables secondary logging should the primary lack sufficient available bandwidth. This mechanism is best used to send an alert to a completely separate destination.

log-response-by-default

Indicates if response logging may be overridden via iRule. This field determines the default response action.

log-response-logging-errors

Enables secondary logging should the primary lack sufficient available bandwidth. This mechanism is best used to send an alert to a completely separate destination.

partition

Displays the administrative partition within which the profile resides.

proxy-close-on-error

Specifies, if enabled, that the logging profile will close the connection after sending its proxy-response.

proxy-respond-on-logging-error

Specifies that the logging profile respond directly (for example, with an HTTP 502) if the logging fails.

proxy-response

Specifies the response to send on logging errors.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

request-log-error-pool

Specifies the name of the pool from which to select log servers.

request-log-error-protocol

Specifies the HighSpeedLogging protocol to use when logging.

request-log-error-template

Specifies the template to use when generating log messages. Shell style escapes (for example, \$foo and/or \${foo}) are used to import transaction-specific values.

request-log-pool

Specifies the name of the pool from which to select log servers.

request-log-protocol

Specifies the HighSpeedLogging protocol to use when logging.

request-log-template

Specifies the template to use when generating log messages. Shell style escapes (for example, \$foo and/or \${foo}) are used to import transaction-specific values.

request-logging

Enables or disables logging before the response is returned to the client.

response-log-error-pool

Specifies the name of the pool from which to select log servers.

response-log-error-protocol

Specifies the HighSpeedLogging protocol to use when logging.

response-log-error-template

Specifies the template to use when generating log messages. Shell style escapes (for example, \$foo and/or \${foo}) are used to import transaction-specific values.

response-log-pool

Specifies the name of the pool from which to select log servers.

response-log-protocol

Specifies the HighSpeedLogging protocol to use when logging.

response-log-template

Specifies the template to use when generating log messages. Shell style escapes (for example, \$foo and/or \${foo}) are used to import transaction-specific values.

response-logging

Enables or disables logging before the response is returned to the client.

to-folder

request-log profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, ltm profile, ltm virtual, modify, mv, show, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 ltm profile request-log(1)

Itm profile response-adapt

NAME

response-adapt - Configures a HTTP response adaptation profile.

MODULE

ltm profile

SYNTAX

Configure the response-adapt component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create response-adapt [name]

modify response-adapt [name]

options:

defaults-from [[name] | none]

enabled [yes | no]

internal-virtual [[name] | none]

preview-size [integer]

service-down-action [ignore | reset | drop]

timeout [integer]

allow-http-10 [yes | no]

edit response-adapt [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

mv response-adapt [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats response-adapt

reset-stats response-adapt [[[name] | [glob] | [regex]] ...]

DISPLAY

list response-adapt

list response-adapt [[[name] | [glob] | [regex]] ...]

show running-config response-adapt

show running-config response-adapt [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

show response-adapt

show response-adapt [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE

delete response-adapt [name]

DESCRIPTION

You can use the response-adapt component to manage a HTTP response adaptation profile.

EXAMPLES

```
create response-adapt my_req_adapt defaults-from response-adapt
```

Creates a HTTP response adaptation profile named my_req_adapt using the system defaults.

```
create response-adapt my_req_adapt { enabled yes }
```

Creates a HTTP response adaptation profile named my_req_adapt that is enabled for adapting HTTP responses.

```
mv response-adapt /Common/my_responseadapt_profile to-folder /Common/my_folder
```

Moves a custom responseadapt profile named my_responseadapt_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is responseadapt.

enabled

Enables adaptation of HTTP responses. If set to yes, HTTP responses will be forwarded to the specified internal virtual server for adaptation. The default value is yes.

internal-virtual

Specifies the name of the internal virtual server to use for adapting the HTTP response.

preview-size

Specifies the maximum size of the preview buffer. The preview buffer is used to hold a copy of the HTTP response header and data sent to the internal virtual server in case the adaptation server reports that it does not need to adapt the HTTP response. Setting the preview-size to 0, disables buffering the response and should only be done if the adaptation server will always return with a modified HTTP response or the original HTTP response. The default value is 1024.

service-down-action

Specifies the action to take if the internal virtual server does not exist or returns an error. The default value is ignore.

The options are:

ignore

Ignore the error and send the unmodified HTTP response to a HTTP server selected from this virtual server's pool.

drop Drop the connection.

reset

Reset the connection.

timeout

Specifies a timeout in milliseconds. If the internal virtual server does not return a result within the specified time, a timeout error will occur. If preview-size is 0, the countdown will start after the entire content has been sent for adaptation. A 0 value disables the timeout. The default value is 0.

to-folder

response-adapt profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

allow-http-10

Specifies whether to forward HTTP version 1.0 responses for adaptation. By default only HTTP version 1.1 responses are forwarded. Version 1.0 is not supported. While it should work in most cases, it might be necessary to restrict adaptation on a site-specific basis. The default value is no.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh, ltm profile request-adapt

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2017-05-24 ltm profile response-adapt(1)

ltm profile rewrite

NAME

rewrite - configure a rewrite profile

MODULE

ltm profile

SYNTAX

Configure the rewrite component within the profile module using the syntax shown in the following sections.

DISPLAY

list rewrite

list rewrite [[name] | [glob]]

show running-config rewrite

show running-config rewrite [[name] | [glob]]

options:

all-properties

non-default-properties

one-line

| grep

show rewrite

show rewrite [[[name] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

CREATE/MODIFY

create rewrite [name]

modify rewrite [name]

options:

app-service [[string] | none]

bypass-list [add | delete | replace-all-with | none] { [uri list] }

client-caching-type [cache-all | cache-css-js | cache-img-css-js | no-cache]

defaults-from [[name] | none]

java-ca-file [[certificate file] | none]

java-crl [[certificate revocation list file] | none]

java-sign-key [[certificate key file] | none]

java-sign-key-passphrase [[string] | none]

java-signer [[certificate file] | none]

location-specific [false | true]

rewrite-list [add | delete | replace-all-with | none] { [uri list] }

rewrite-mode [portal | uri-translation]

set-cookie-rules [add | delete | modify | replace-all-with | none] {

[name] {

client {

domain [string]

path [string]

}

server {

domain [string]

path [string]

}

}

}

```

split-tunneling [false | true]
uri-rules [add | delete | modify | replace-all-with | none] {
  [name] {
[type [both | request | response]]
client {
  scheme [string]
  host [string]
  port [string]
  path [string]
}
server {
  scheme [string]
  host [string]
  port [string]
  path [string]
}
}
}

```

```

edit rewrite [ [ [name] | [glob] ] ... ]
options:
  all-properties
  non-default-properties

```

```

DELETE
delete rewrite [name]

```

DESCRIPTION

Use the rewrite component to configure a Rewrite Profile in URI Translation or Portal (Access) mode.

EXAMPLES

URI Translation Mode

Create a profile

```
create my_uri_rewrite rewrite-mode uri-translation
```

Add a rule to rewrite URIs

```
modify my_uri_rewrite uri-rules add { my_rule { client { path /client/ } server { path /server/ } } }
```

```
modify my_uri_rewrite uri-rules add { my_rule { client { scheme http host www.client.com path / }
server { scheme http host www.server.com path / } } }
```

Add a rule to rewrite Set-Cookie headers

```
modify my_uri_rewrite set-cookie-rules add { my_rule { client { domain client.com path / } server {
domain server.com path / } } }
```

Portal (Access) Mode

Create a profile

```
create my_portal_rewrite rewrite-mode portal
```

Configure the client to cache all files

```
modify my_portal_rewrite client-caching-type cache-all
```

Set the rewrite list and bypass list

```
modify my_portal_rewrite rewrite-list add { *://www.myportal.com/* http://abc*.com/* } bypass-list add
{ *://external_web.com/* }
```

Configure split-tunneling

```
modify my_portal_rewrite split-tunneling true
```

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

bypass-list

Specifies a list of URIs that are bypassed inside a web page when the page is accessed using Portal Access. The default is none.

client-caching-type

Specifies one of four options for client caching. When the Client Cache setting for a web application resource is set to default, the system uses the setting configured in the Rewrite profile. If the Client Cache option is configured for any other setting, the web application resource item caching configuration overwrites the setting in the Rewrite profile. The default is cache-css-js. The options are:

cache-all

Do not modify cache headers on backend servers.

cache-css-js

Cache only the CSS file and Java Script.

cache-img-css-js

Cache only images, the CSS file and Java Script.

no-cache

Eliminate caching.

defaults-from

Specifies the profile from which the Rewrite profile inherits properties. Explicitly specified properties override inherited properties.

java-ca-file

Specifies a CA against which to verify signed Java applets signatures. The default value is ca-bundle.crt.

java-crl

Specifies a CRL against which to verify signed Java applets signature certificates. The default value is none.

java-sign-key

Specifies a private key for re-signing of signed Java applets after patching. The default value is default.key.

java-sign-key-passphrase

Specifies a passphrase for the private key to be encrypted with. The default value is none. Note: your passphrase will be encrypted and displayed under the label java-sign-key-passphrase-encrypted.

java-signer

Specifies a certificate to use for re-signing of signed Java applets after patching. The default value is default.crt.

location-specific

Specifies whether or not this object contains one or more attributes with values that are specific to the location where the BIG-IP device resides. The location-specific attribute is either true or false. When using policy sync, mark an object as location-specific to prevent errors that can occur when policies reference objects, such as authentication servers, that are specific to a certain location. The default value is none.

rewrite-list

Specifies a list of URIs that are rewritten inside a web page when the page is accessed using Portal Access. The default value is none.

rewrite-mode

Specifies the mode of rewriting. uri-translation is a rules-based rewrite mode. portal is for use with Portal Access.

set-cookie-rules

Used with uri-translation mode. Specifies the rules for rewriting HTTP Set-Cookie headers. Each rule has a name and a client and server domain and path. The name may be any alphanumeric string and must be unique. The path must be an absolute directory path and not a relative path or a file path. If the domain and path of the Set-Cookie header in the HTTP response match the domain and path of the server side of a rule, they will be rewritten to the domain and path of client side of that rule. Set-Cookie rules take precedence over URI rules when rewriting Set-Cookie headers.

split-tunneling

Specifies whether the profile provides for split tunneling. The default is false.

uri-rules

Used with uri-translation mode. Specifies the rules for rewriting request and response headers and response bodies. These rules affect the following.

request headers

URI, Host, Referer

response headers

Content-Location, Link, Location, Refresh, Set-Cookie

response body

HTML, CSS

Each rule has a name, a type, and a client and server URI. The name may be any alphanumeric string and must be unique. The type may be "request", "response", or "both": "request" rules affect request headers only, "response" rules affect response headers and bodies only, and "both" rules affect both. URIs must include a path; scheme, host, and port are optional. If a URI must contain a scheme or host, it must include both. If it must include a port, it must also include a scheme and host. Paths may be absolute directory paths only. They may not be relative paths or file paths. If a URI in a request header matches the client side URI of a rule, it will be rewritten to the server side URI of that rule. If a URI in a response header or body matches the server side URI of a rule, it will be rewritten to the client side URI of that rule. When rewriting Set-Cookie headers, the host and path of the server side URI are used to match the domain and path of the header. The client side host and path replace that header's domain and path if a match is found. Set-Cookie rules take precedence over URI rules when rewriting Set-Cookie headers.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015. All rights reserved.

Itm profile rtsp

NAME

rtsp - Configures an RTSP (realtime streaming protocol) profile.

MODULE

itm profile

SYNTAX

Configure the rtsp component within the itm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create rtsp [name]

modify rtsp [name]

options:

app-service [[string] | none]

check-source [disabled | enabled]

defaults-from [name]

description [string]

idle-timeout [integer]

max-header-size [integer]

max-queued-data [integer]

multicast-redirect [disabled | enabled]

proxy [external | internal | none]

proxy-header [[name] | none]

real-http-persistence [disabled | enabled]

rtcp-port [number]

rtp-port [number]

session-reconnect [disabled | enabled]

unicast-redirect [disabled | enabled]

log-publisher [log publisher name | none]

log-profile [log profile name | none]

edit rtsp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv rtsp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats rtsp

reset-stats rtsp [[[name] | [glob] | [regex]] ...]

DISPLAY

list rtsp

list rtsp [[[name] | [glob] | [regex]] ...]

show running-config rtsp

show running-config rtsp

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show rtsp

show rtsp [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete rtsp [name]

DESCRIPTION

You can use the rtsp component to manage a profile that you use to control RTSP traffic.

EXAMPLES

```
create rtsp my_rtsp_profile defaults-from rtsp
```

Creates a custom RTSP profile named my_rtsp_profile that inherits its settings from the system default RTSP profile.

```
list rtsp all-properties
```

Displays all properties for all RTSP profiles.

```
mv rtsp /Common/my_rtsp_profile to-folder /Common/my_folder
```

Moves a custom rtsp profile named my_rtsp_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

check-source

When enabled the system uses the source attribute in the transport header to establish the target address of the RTP stream, and before the response is forwarded to the client, updates the value of the source attribute to be the virtual address of the BIG-IP system. When disabled the system does not change the source attribute. The default value is enabled.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is rtsp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

idle-timeout

Specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 300 seconds.

max-header-size

Specifies the maximum size of an RTSP request or response header that the RTSP filter accepts before dropping the connection. The default value is 4096 bytes.

max-queued-data

Specifies the maximum amount of data that the RTSP filter buffers before dropping the connection. The default value is 32768 bytes.

multicast-redirect

Specifies whether to enable or disable multicast redirect. When enabled, the client can select the destination to which to stream data. The default value is disabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

proxy

Specifies whether the RTSP filter is associated with an RTSP proxy configuration. The default value is none.

proxy-header

When the proxy option is set, specifies the name of the header in the RTSP proxy configuration that is passed from the client-side virtual server to the server-side virtual server. Note that the name of the header must begin with X-. The default value is none.

To use the proxy-header option, you must specify a value for the proxy option.

real-http-persistence

Specifies whether to enable or disable real HTTP persistence. When enabled, the RTSP filter automatically persists Real Networks RTSP over HTTP using the RTSP port. The default value is enabled. If you disable this parameter, you can override the default behavior with an iRule.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rtcp-port

Specifies the number of the port to use for the Real Time Control Protocol (RTCP) service. The default value is 0 (zero). RTCP allows monitoring of real-time data delivery.

rtp-port

Specifies the number of the port to use for the RTP service. The default value is 0 (zero).

session-reconnect

Specifies whether to enable or disable session reconnect. When enabled, the RTSP filter persists the control connection, which is being resumed, to the correct server. The default value is disabled.

to-folder

rtsp profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

unicast-redirect

Specifies whether to enable or disable unicast redirect. When enabled, the client can select the destination to which to stream data. The default value is disabled.

log-publisher

Specify the name of the log publisher which logs translation events. See help sys log-config for more details on the logging sub-system. Use the "sys log-config publisher" component to set up a log publisher.

log-profile

Specify the name of the ALG log profile which controls the logging of ALG . See help ltm alg-log-profile for more details on the logging profile sub-system. Use the "ltm alg-log-profile profile" component to set up a ALG log profile.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 ltm profile rtsp(1)

ltm profile sctp

NAME

sctp - Configures a Stream Control Transmission Protocol (SCTP) profile.

MODULE

ltm profile

SYNTAX

Configure the sctp component within the ltm profile module using the syntax shown in the following sections.

CREATE

create sctp [name]

modify sctp [name]

options:

app-service [[string] | none]

cookie-expiration [integer]

defaults-from [name]

description [string]

heartbeat-interval [integer]

heartbeat-max-burst [integer]

idle-timeout [integer]

in-streams [integer]

init-max-retries [integer]

ip-tos [integer]

link-qos [integer]

max-burst [integer]

out-streams [integer]

proxy-buffer-high [integer]

proxy-buffer-low [integer]

receive-chunks [integer]

receive-ordered [disabled | enabled]

receive-window-size [integer]

reset-on-timeout [disabled | enabled]

rto-initial [integer]

rto-max [integer]

rto-min [integer]

sack-timeout [integer]

secret [default | [string]]

send-buffer-size [integer]

send-max-retries [integer]

send-partial [disabled | enabled]

tcp-shutdown [disabled | enabled]

transmit-chunks [integer]

edit sctp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv sctp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats sctp
reset-stats sctp [[name] | [glob] | [regex]] ...]

DISPLAY
list sctp
list sctp [[name] | [glob] | [regex]] ...]
show running-config sctp
show running-config sctp
[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

show sctp
show sctp [[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete sctp [name]

DESCRIPTION
You can use the sctp component to manage a profile for SCTP traffic.

EXAMPLES
create sctp my_sctp_profile defaults-from sctp

Creates a custom SCTP profile named my_sctp_profile that inherits its settings from the system default SCTP profile.

list sctp all-properties

Displays all properties for all SCTP profiles.

mv sctp /Common/my_sctp_profile to-folder /Common/my_folder

Moves a custom sctp profile named my_sctp_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS
app-service
Specifies the name of the application service to which the profile belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

cookie-expiration
Specifies how many seconds the cookie is valid. The default value is 60 seconds.

defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is sctp.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

heartbeat-interval
Specifies the number of seconds to wait before sending a heartbeat chunk. The default value is 30 seconds.

heartbeat-max-burst
Specifies the number of heartbeat packets to be sent in a single burst. The default value is 1.

idle-timeout
Specifies the number of seconds without traffic before a connection is eligible for deletion. The default value is 300 seconds.

in-streams
Specifies the number of inbound streams to advertise on new connections. The default value is 1.

init-max-retries
Specifies the maximum number of retries to establish a connection. The default value is 4.

ip-tos
Specifies the Type of Service (ToS) that is set in packets sent to the peer. The default value is 0.

link-qos
Specifies the Link Quality of Service (QoS) that is set in sent packets. The default value is 0.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`out-streams`

Specifies the number of outbound streams. The default value is 2.

`partition`

Displays the administrative partition within which the component resides.

`proxy-buffer-high`

Specifies the proxy buffer level after which the system closes the receive window. The default value is 16384.

`proxy-buffer-low`

Specifies the proxy buffer level after which the system opens the receive window. The default value is 4096.

`receive-chunks`

Specifies the size (in chunks) of the `rx_chunk` buffer. The default value is 256.

`receive-ordered`

When enabled, the default, the system delivers messages to the application layer in order.

`receive-window-size`

Specifies the size (in bytes) of the receive window. Prorate this value to the `receive-chunks` value. The default value is 65535.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`reset-on-timeout`

When enabled, the default, the system resets the connection when the connection times out.

`rto-initial`

Specifies the number of milliseconds for the initial value of retransmission timeout. The default value is 3000 milliseconds.

`rto-max`

Specifies the number of milliseconds for the maximum value of retransmission timeout. The default value is 60000 milliseconds.

`rto-min`

Specifies the number of milliseconds for the minimum value of retransmission timeout. The default value is 1000 milliseconds.

`sack-timeout`

Specifies the number of milliseconds for the delayed selective acknowledgement timeout. The default value is 200 milliseconds.

`secret`

Specifies the internal secret string used for HTTP Message Authenticated Code (HMAC) cookies.

`send-buffer-size`

Specifies the size in bytes of the buffer. The default value is 65536.

`max-burst`

Specifies the maximum number of data packets to send in a single burst. The default value is 4.

`send-max-retries`

Specifies the maximum number of time the system tries again to send the data. The default value is 8.

`send-partial`

When enabled, the default, the system accepts partial application data.

`tcp-shutdown`

When enabled, the system emulates the closing of a TCP connection. The default value is enabled.

`to-folder`

sctp profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

`transmit-chunks`

Specifies the size of the `tx_chunk` buffer. The default value is 256.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `mv`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015. All rights reserved.

Itm profile server-ldap

NAME

server-ldap - Configures an Server LDAP profile.

MODULE

itm profile

SYNTAX

Configure the server-ldap component within the itm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create server-ldap [name]

modify server-ldap [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

activation-mode [none | allow | require]

ss-activation-mode [none | allow | require]

edit server-ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list server-ldap

list server-ldap [[[name] | [glob] | [regex]] ...]

show running-config server-ldap

show running-config server-ldap [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DELETE

delete server-ldap [name]

DESCRIPTION

You can use the server-ldap component to create, modify, display, or delete an Server LDAP profile with which you can manage Server LDAP traffic.

EXAMPLES

```
create server-ldap my_serverldap_profile defaults-from serverldap
```

Creates a custom Server LDAP profile named my_serverldap_profile that inherits its settings from the system default Server LDAP profile.

```
list server-ldap
```

Displays the properties of all Server LDAP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is smtp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

activation-mode

Sets the activation-mode for STARTTLS. The options are NONE, ALLOW, or REQUIRE. The default value is

NONE.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, sys provision, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-05-06 ltm profile server-ldap(1)

ltm profile server-ssl

NAME

server-ssl - Configures a Server SSL profile.

MODULE

ltm profile

SYNTAX

Configure the server-ssl component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create server-ssl [name]

modify server-ssl [name]

options:

alert-timeout [indefinite | immediate | [integer]]
allow-expired-crl [enabled | disabled]
app-service [[string] | none]
authenticate [always | once]
authenticate-depth [integer]
authenticate-name [[name] | none]
bypass-on-client-cert-fail [disabled | enabled]
bypass-on-handshake-alert [disabled | enabled]
c3d-ca-cert [name]
c3d-ca-key [name]
c3d-ca-passphrase [string]
c3d-cert-extension-custom-oids [none | [string]]
c3d-cert-extension-includes {
none |
[basic-constraints extended-key-usage
key-usage subject-alternative-name
]...
}
c3d-cert-lifespan [integer]
ca-file [[file name] | none]
cache-size [integer]
cache-timeout [integer]
cert [[file name] | none]
chain [[name] | none]
cipher-group [name | none]
ciphers [[name] | none]
crl [[name] | none]
crl-file [none]
defaults-from [[name] | none]
description [string]
expire-cert-response-control [drop | ignore | mask]
handshake-timeout [indefinite | [integer]]
key [[file name] | none]
max-active-handshakes [integer]
mod-ssl-methods [disabled | enabled]
mode [disabled | enabled]
ocsp [[ocsp profile name] | none]
options {
none |
[dont-insert-empty-fragments
no-session-resumption-on-renegotiation
no-ssl no-sslv3 no-tls no-tlsv1 no-tlsv1.1 no-tlsv1.2
no-tlsv1.3 no-dtls no-dtlsv1.0 no-dtlsv1.2 gmsslv1.1 passive-close
single-dh-use tls-rollback-bug]
}
passphrase [none | [string]]
peer-cert-mode [ignore | require]
proxy-ssl [disabled | enabled]
proxy-ssl-passthrough [disabled | enabled]

renegotiate-period [indefinite | [integer]]
renegotiate-size [indefinite | [integer]]
renegotiation [disabled | enabled]
retain-certificate [true | false]
revoked-cert-status-response-control [drop | ignore | mask]
secure-renegotiation [request | require | require-strict]
server-name [name]
session-mirroring [disabled | enabled]
session-ticket [disabled | enabled]
generic-alert [disabled | enabled]
sni-default [true | false]
sni-require [true | false]
ssl-c3d [disabled | enabled]
ssl-forward-proxy [disabled | enabled]
ssl-forward-proxy-bypass [disabled | enabled]
ssl-forward-proxy-verified-handshake [disabled | enabled]
ssl-sign-hash [any | sha1 | sha256 | sha384]
strict-resume [disabled | enabled]
unclean-shutdown [disabled | enabled]
data-Ortt [disabled | enabled]
unknown-cert-status-response-control [ignore | drop | mask]
untrusted-cert-response-control [drop | ignore | mask]

edit server-ssl [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv server-ssl [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats server-ssl

reset-stats server-ssl [[[name] | [glob] | [regex]] ...]

DISPLAY

list server-ssl

list server-ssl [[[name] | [glob] | [regex]] ...]

show running-config server-ssl

show running-config server-ssl

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show server-ssl

show server-ssl [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

global

DELETE

delete server-ssl [all | [name]]

options:

recursive

DESCRIPTION

You can use the server-ssl component to manage a server SSL profile.

Server-side profiles enable the traffic management system to handle encryption tasks for any SSL connection being sent from a local traffic management system to a target server. A server-side SSL profile acts as a client by presenting certificate credentials to a server when authentication of the local traffic management system is required. You implement this type of profile by using the default profile, or by creating a custom profile based on the Server SSL profile template and modifying its settings.

EXAMPLES

```
create server-ssl my_serverssl_profile defaults-from serverssl
```

Creates a custom Server SSL profile named my_serverssl_profile that inherits its settings from the system default profile serverssl.

```
list server-ssl all-properties
```

Displays all properties for all Server SSL profiles.

```
mv server-ssl /Common/my_serverssl_profile to-folder /Common/my_folder
```

Moves a custom server-ssl profile named my_serverssl_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

alert-timeout

Specifies the maximum time period in seconds to keep the SSL session active after alert message is sent, or indefinite. The default value is indefinite.

allow-expired-crl

Use the specified CRL file even if it has expired. The default value is disabled.

authenticate

Specifies the frequency of authentication. The default value is once. Note that if this is set to always session cache and session ticket will be disabled.

authenticate-depth

Specifies the client certificate chain maximum traversal depth. The default value is 9.

authenticate-name

Specifies a Common Name (CN) that is embedded in a server certificate. The system authenticates a server based on the specified CN. The default value is none.

bypass-on-client-cert-fail

Enables or disables SSL forward proxy bypass on failing to get client certificate that server asks for. When enabled and the SSL handshake cannot be completed because of failure to get the client certificate, SSL traffic bypasses the BIG-IP system untouched, without decryption/encryption. The default value is disabled. Conversely, you can specify enabled to use this feature.

bypass-on-handshake-alert

Enables or disables SSL forward proxy bypass on receiving handshake_failure, protocol_version or unsupported_extension alert message during the serverside SSL handshake. When enabled and there is an SSL handshake_failure, protocol_version or unsupported_extension alert during the serverside SSL handshake, SSL traffic bypasses the BIG-IP system untouched, without decryption/encryption. The default value is disabled. Conversely, you can specify enabled to use this feature.

c3d-ca-cert

Specifies the name of the certificate file that is used as the certification authority certificate when SSL client certificate constrained delegation is enabled. The certificate should be generated and installed by you on the system. When selecting this option, type a certificate file name.

c3d-ca-key

Specifies the name of the key file that is used as the certification authority key when SSL client certificate constrained delegation is enabled. The key should be generated and installed by you on the system. When selecting this option, type a key file name.

c3d-ca-passphrase

Specifies the passphrase of the key file that is used as the certification authority key when SSL client certificate constrained delegation is enabled. When selecting this option, type the passphrase corresponding to the selected c3d-ca-key.

c3d-cert-extension-custom-oids

Specifies the custom extension OID of the client certificates to be included in the generated certificates using SSL client certificate constrained delegation.

c3d-cert-extension-includes

Specifies the extensions of the client certificates to be included in the generated certificates using SSL client certificate constrained delegation. For example, { basic-constraints }. The default value is { basic-constraints extended-key-usage key-usage subject-alternative-name }. The extensions are:

basic-constraints

Basic constraints are used to indicate whether the certificate belongs to a CA.

extended-key-usage

Extended Key Usage is used, typically on a leaf certificate, to indicate the purpose of the public key contained in the certificate.

key-usage

Key Usage provides a bitmap specifying the cryptographic operations which may be performed using the public key contained in the certificate; for example, it could indicate that the key should be used for signature but not for encipherment.

subject-alternative-name

Subject Alternative Name allows identities to be bound to the subject of the certificate. These identities may be included in addition to or in place of the identity in the subject field of the certificate.

c3d-cert-lifespan

Specifies the lifespan of the certificate generated using the SSL client certificate constrained delegation. The default value is 24.

ca-file

Specifies the certificate authority file name. Configures certificate verification by specifying a list of client or server CAs that the traffic management system trusts. The default value is none.

cache-size

Specifies the SSL session cache size. For client profiles only, you can configure timeout and size values for the SSL session cache. Because each profile maintains a separate SSL session cache, you can configure the values on a per-profile basis. The default value is 262144.

cache-timeout

Specifies the SSL session cache timeout value, which is the usable lifetime seconds of negotiated SSL session IDs. The default value is 3600 seconds. Acceptable values are integers greater than or equal to 0 and less than or equal to 86400.

`cert` Specifies the name of the certificate installed on the traffic management system for the purpose of terminating or initiating an SSL connection. The default value is none.

`chain`
Specifies or builds a certificate chain file that a client can use to authenticate the profile. The default value is none.

`cipher-group`
Specifies a cipher group. If the cipher group is not blank or none, the ciphers string will be used.

`ciphers`
Specifies a cipher name. The default value is DEFAULT.

`crl` Specifies the name of crl validator for validating status of server certificate. Specifying none disables crl validation of server certificate. The default value is none.

`crl-file`
Specifies the certificate revocation list file name. The default value is none.

`defaults-from`
Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is serverssl.

`description`
User defined description.

`expire-cert-response-control`
Specifies the BIGIP action when the server certificate has expired. The default value is drop, which causes the connection to be dropped. Conversely, you can specify ignore to cause the connection to ignore the error and continue or you can specify mask in case of SSL forward proxy to mask server certificate errors and continue with handshake and forge a good certificate on client-side.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`handshake-timeout`
Specifies the handshake timeout in seconds. The default value is 10.

`key` Specifies the key file name. Specifies the name of the key installed on the traffic management system for the purpose of terminating or initiating an SSL connection. The default value is none.

`mod-ssl-methods`
Enables or disables ModSSL methods. The default value is disabled.

Enable this option when OpenSSL methods are inadequate. For example, you can enable ModSSL method emulation when you want to use SSL compression over TLSv1.

`mode` Enables or disables SSL processing. The default value is enabled.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`ocsp` Specifies the name of ocsp profile for purpose of validating status of server certificate. Specifying none disables ocsp validation of server certificate. The default value is none.

`options`
Enables options, including some industry-related workarounds. Enter options inside braces, for example, {dont-insert-empty-fragments}. The default value is dont-insert-empty-fragments no-tlsv1.3.

`dont-insert-empty-fragments`
Disables a countermeasure against a SSL 3.0/TLS 1.0 protocol vulnerability affecting CBC ciphers. These ciphers cannot be handled by certain broken SSL implementations. This option has no effect for connections using other ciphers.

`max-active-handshakes`
Specifies the maximum number allowed SSL active handshakes. The default value is 0.

`no-session-resumption-on-renegotiation`
When performing renegotiation as an SSL server, this option always starts a new session (that is, session resumption requests are accepted only in the initial handshake). The system ignores this option for server-side SSL.

`gmssl1.1`
Enable GMSSLv1.1 protocol.

`no-ssl`
Do not use any version of the SSL protocol.

`no-sslv3`
Do not use the SSLv3 protocol.

`no-tls`
Do not use any version of the TLS protocol.

`no-tls1`
Do not use the TLSv1.0 protocol.

`no-tls1.1`
Do not use the TLSv1.1 protocol.

`no-tls1.2`
Do not use the TLSv1.2 protocol.

`no-tls1.3`
Do not use the TLSv1.3 protocol. Note that this is for future expansion. Currently TLSv1.3 has not been implemented for server side SSL, so removing this will have no effect and log a warning message.

`no-dtls`
Do not use any version of the DTLS protocol.

`no-dtls1.0`
Do not use the DTLSv1.0 protocol.

`no-dtls1.2`
Do not use the DTLSv1.2 protocol.

`passive-close`
Specifies how to handle passive closes.

`none` Disables all workarounds. Note that F5 Networks does not recommend this option.

`single-dh-use`
Creates a new key when using temporary/ephemeral DH parameters. This option must be used to prevent small subgroup attacks, when the DH parameters were not generated using strong primes (for example, when using DSA-parameters). If strong primes were used, it is not strictly necessary to generate a new DH key during each handshake, but F5 Networks recommends it. Enable the Single DH Use option whenever temporary or ephemeral DH parameters are used.

`tls-rollback-bug`
Disables version rollback attack detection. During the client key exchange, the client must send the same information about acceptable SSL/TLS protocol levels as it sends during the first hello. Some clients violate this rule by adapting to the server's answer. For example, the client sends an SSLv2 hello and accepts up to SSLv3.1 (TLSv1), but the server only processes up to SSLv3. In this case, the client must still use the same SSLv3.1 (TLSv1) announcement. Some clients step down to SSLv3 with respect to the server's answer and violate the version rollback protection. The system ignores this option for server-side SSL.

`partition`
Displays the administrative partition within which the component resides.

`passphrase`
Specifies the key passphrase, if required. The default value is none.

`peer-cert-mode`
Specifies the peer certificate mode. The default value is ignore.

`proxy-ssl`
Enabling this option requires a corresponding client ssl profile with proxy-ssl enabled to perform transparent SSL decryption. This feature allows further modification of application traffic within an SSL tunnel while still allowing the server to perform necessary authorization, authentication, auditing steps.

`proxy-ssl-passthrough`
Enabling this option requires a corresponding client ssl profile with proxy-ssl-passthrough enabled. This allows Proxy SSL to passthrough the traffic when ciphersuite negotiated between the client and server is not supported. The default option is disabled.

`regex`
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`renegotiate-period`
Specifies the number of seconds from the initial connect time after which the system renegotiates an SSL session. The default value is indefinite, which means that you do not want the system to renegotiate SSL sessions.

Each time the session renegotiation is successful, a new connection is started. Therefore, the system attempts to renegotiate the session again, in the specified amount of time following a successful session renegotiation. For example, setting the renegotiate-period option to 3600 seconds triggers session renegotiation at least once an hour.

`renegotiate-size`
Specifies a throughput size, in megabytes, of SSL renegotiation. This option forces the traffic management system to renegotiate an SSL session based on the size, in megabytes, of application data that is transmitted over the secure channel. The default value is indefinite, which specifies that you do not want a throughput size.

`renegotiation`

Specifies whether renegotiations are enabled. The default value is enabled. When renegotiations are disabled, the system is acting as an SSL server, and a COMPAT or NATIVE cipher is negotiated, the system will abort the connection. Additionally, when renegotiations are disabled and the system is acting as an SSL client, the system will ignore the server's HelloRequest messages.

retain-certificate

APM module requires storing certificate in SSL session. When set to false, certificate will not be stored in SSL session. The default value is true.

revoked-cert-status-response-control

Specifies the BIGIP action when the server certificate status is revoked. The default value is drop, which causes the connection to be dropped. You can specify ignore to cause the connection to ignore the error and continue handshake. You can specify mask in case of SSL forward proxy to mask server certificate status error and continue handshake.

generic-alert

Enables or disables generic-alert. The default option is enabled, which causes the SSL profile to use generic alert number. Conversely, you can specify disabled to cause SSL profile to use alert number defined in RFC5246/RFC6066 strictly.

secure-renegotiation

Specifies the secure renegotiation mode. The default value is require-strict. When secure renegotiation is set to require, any connection to an unpatched server will be aborted. For server-ssl, there is no difference between require and require-strict secure renegotiation. When secure renegotiation is set to request, connections to unpatched servers will be permitted. This setting is NOT recommended however, as it is subject to active man-in-the-middle attacks.

server-name

Specifies the server name to be included in SNI (server name indication) extension during SSL handshake in ClientHello.

session-mirroring

Enables or disables the mirroring of sessions to high availability peer. By default, this setting is disabled, which causes the system to not mirror ssl sessions.

session-ticket

Enables or disables session-ticket. The default option is disabled, which causes the SSL profile not to use session ticket per RFC 5077. Conversely, you can specify enabled to cause SSL profile to use session ticket per RFC 5077.

sni-default

When true, this profile is the default SSL profile when the server name in a client connection does not match any configured server names, or a client connection does not specify any server name at all.

sni-require

When this option is enabled, connections to a server that does not support SNI extension will be rejected.

ssl-c3d

Enables or disables SSL Client certificate constrained delegation. The default option is disabled. Conversely, you can specify enabled to use the SSL client certificate constrained delegation.

ssl-forward-proxy

Enables or disables ssl-forward-proxy feature. The default option is disabled. Conversely, you can specify enabled to use the SSL Forward Proxy Feature.

ssl-sign-hash

Specifies SSL sign hash algorithm which is used to sign and verify SSL Server Key Exchange and Certificate Verify messages for the specified SSL profiles. The default value is sha1.

ssl-forward-proxy-bypass

Enables or disables ssl-forward-proxy-bypass feature. The default option is disabled. Conversely, you can specify enabled to use the SSL Forward Proxy Bypass Feature.

ssl-forward-proxy-verified-handshake

Specifies, when enabled, that in SSL forward proxy mode, the system should always do a TLS handshake with the server first before doing the client handshake. When disabled, the system will do the server handshake first only if it has not previously forged and cached the server certificate; once the server certificate is ready, the system will always handshake first with the client. The default value is disabled.

strict-resume

Enables or disables the resumption of SSL sessions after an unclean shutdown. The default value is disabled, which indicates that the SSL profile refuses to resume SSL sessions after an unclean shutdown.

to-folder

server-ssl profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

unclean-shutdown

Specifies, when enabled, that the SSL profile performs unclean shutdowns of all SSL connections, which means that underlying TCP connections are closed without exchanging the required SSL shutdown alerts. If you want to force the SSL profile to perform a clean shutdown of all SSL connections, you can disable this option.

unknown-cert-status-response-control

Specifies the BIGIP action when the server certificate status is unknown. The default value is ignore,

which causes the connection to ignore the error and continue handshake. You can specify drop which causes the connection to be dropped. You can specify mask in case of SSL forward proxy to mask server certificate status error and continue handshake.

untrusted-cert-response-control

Specifies the BIGIP action when the server certificate has untrusted CA. The default value is drop, which causes the connection to be dropped. Conversely, you can specify ignore to cause the connection to ignore the error and continue or you can specify mask in case of SSL forward proxy to mask server certificate errors and continue with handshake and forge a good certificate on client-side.

data-0rtt

Specifies if TLSv1.3 should send 0-RTT early data when available. The default value is disabled.

SEE ALSO

create, delete, edit, glob, list, ltm profile client-ssl, ltm virtual, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2020-02-08 ltm profile server-ssl(1)

ltm profile sip

NAME

sip - Configures a Session Initiation Protocol (SIP) profile.

MODULE

ltm profile

SYNTAX

Configure the sip component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create sip [name]

modify sip [name]

options:

alg-enable [disabled | enabled]

app-service [[string] | none]

community [[community name] | none]

defaults-from [[name] | none]

description [string]

dialog-aware [disabled | enabled]

dialog-establishment-timeout [integer]

enable-sip-firewall [no | yes]

insert-record-route-header [disabled | enabled]

insert-via-header [disabled | enabled]

max-media-sessions [integer]

max-registrations [integer]

max-sessions-per-registration [integer]

max-size [integer]

registration-timeout [integer]

rtp-proxy-style [symmetric | restricted-by-ip-address | any-location]

secure-via-header [disabled | enabled]

security [disabled | enabled]

sip-session-timeout [integer]

terminate-on-bye [disabled | enabled]

user-via-header [[via-header] | none]

log-publisher [log publisher name | none]

log-profile [log profile name | none]

edit sip [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv sip [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats sip

reset-stats sip [[[name] | [glob] | [regex]] ...]

DISPLAY

list sip

list sip [[[name] | [glob] | [regex]] ...]

```
show running-config sip
show running-config sip [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show sip
show sip [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

```
DELETE
delete sip [name]
```

DESCRIPTION

You can use the sip component to manage a SIP profile.

EXAMPLES

```
create sip my_sip_profile defaults-from sip
```

Creates a SIP profile named my_sip_profile using the system defaults.

```
create sip my_sip_profile { terminate-by disabled }
```

Creates a SIP profile named my_sip_profile that leaves a connection open following the completion of a BYE transaction.

```
mv sip /Common/my_sip_profile to-folder /Common/my_folder
```

Moves a custom sip profile named my_sip_profile to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

alg-enable

Enables or disables the SIP ALG (Application Level Gateway) feature. The default value is disabled. Note: for a SIP profile with ALG enabled to function correctly, the virtual which uses the profile must have destination and mask set to 0.0.0.0 for IPv4, or :: for IPv6. Additionally, the virtual must have source-address-translation enabled.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

community

Specifies the community to which you want to assign the virtual server that you associate with this profile. The default value is none.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all of the settings and values from the specified parent profile. The default value is sip.

description

User defined description.

dialog-aware

Enables or disables the ability for the system to be aware of unauthorized use of the SIP dialog. The default value is disabled.

dialog-establishment-timeout

Indicates the timeout value for dialog establishment in a sip session. The default value is 10 seconds.

enable-sip-firewall

Indicates whether to enable SIP firewall functionality or not. Default value is no.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

insert-record-route-header

Enables or disables the insertion of a Record-Route header, which indicates the next hop for the following SIP request messages. The default value is disabled.

insert-via-header

Enables or disables the insertion of a Via header, which indicates where the message originated. The response message uses this routing information. The default value is disabled.

max-media-sessions

Indicates the maximum number of SDP media sessions that the BIG-IP system accepts. The default value is 6.

max-registrations

Indicates the maximum number of registrations, the maximum allowable REGISTER messages can be recorded that the BIG-IP system accepts. The default value is 100.

max-sessions-per-registration

Indicates the maximum number of calls or sessions can be made by a user for a single registration that the BIG-IP system accepts. The default value is 50.

max-size

Specifies the maximum SIP message size that the BIG-IP system accepts. The default value is 65535 bytes.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

registration-timeout

Indicates the timeout value for a sip registration. The default value is 3600 seconds.

rtp-proxy-style

Indicates the style in which the RTP will proxy the data. When a dialog is established, the necessary SDP data needs to know where the RTP flows are directed. The default value is symmetric. The options available are:

symmetric

Indicates the use of a bidirectional related flow.

restricted-by-ip-address

Indicates the use of ephemeral listeners to support fixed client IP, listener is restricted to connections coming from a particular source.

any-location

Indicates the use of ephemeral listeners to support wildcard, connections are allowed to come from anyway.

secure-via-header

Enables or disables the insertion of a Secure Via header, which indicates where the message originated. When you are using SSL/TLS (over TCP) to create a secure channel with the server node, use this setting to configure the system to insert a Secure Via header into SIP requests. The default value is disabled.

security

Enables or disables security for the SIP profile. The default value is disabled.

sip-session-timeout

Indicates the timeout value for a sip session. The default value is 300 seconds.

terminate-on-bye

Enables or disables the termination of a connection when a BYE transaction finishes. Use this parameter with UDP connections only, not with TCP connections. The default value is enabled.

to-folder

sip profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

user-via-header

Enables or disables the insertion of a Via header specified by a system administrator. The default value is none.

log-publisher

Specify the name of the log publisher which logs translation events. See help sys log-config for more details on the logging sub-system. Use the "sys log-config publisher" component to set up a log publisher.

log-profile

Specify the name of the ALG log profile which controls the logging of ALG . See help ltm alg-log-profile for more details on the logging profile sub-system. Use the "ltm alg-log-profile profile" component to set up a ALG log profile.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsk

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2014, 2016. All rights reserved.

Itm profile smtp

NAME

smtp - Configures an SMTP profile.

MODULE

itm profile

SYNTAX

Configure the smtp component within the Itm profile module using the syntax shown in the following sections. The smtp profile is available when the asm module is enabled. You enable the asm module via provisioning commands, which are described in help sys provision.

CREATE/MODIFY

create smtp [name]

modify smtp [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

security [disabled | enabled]

edit smtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list smtp

list smtp [[[name] | [glob] | [regex]] ...]

show running-config smtp

show running-config smtp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DELETE

delete smtp [name]

DESCRIPTION

You can use the smtp component to create, modify, display, or delete an SMTP profile with which you can manage SMTP traffic.

EXAMPLES

```
create smtp my_smtp_profile defaults-from smtp
```

Creates a custom SMTP profile named my_smtp_profile that inherits its settings from the system default SMTP profile.

```
list smtp
```

Displays the properties of all SMTP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is smtp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

security

Enables or disables secure SMTP traffic for the BIG-IP(r) Application Security Manager. The default value is disabled.

SEE ALSO

create, delete, edit, glob, list, Itm virtual, modify, regex, reset-stats, show, sys provision, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2013-11-06 ltm profile smtp(1)

ltm profile smtps

NAME

smtps - Configures an SMTPs profile.

MODULE

ltm profile

SYNTAX

Configure the smtps component within the ltm profile module using the syntax shown in the following sections. The smtps profile is available when the asm module is enabled. You enable the asm module via provisioning commands, which are described in help sys provision.

CREATE/MODIFY

create smtps [name]

modify smtps [name]

options:

app-service [(string) | none]

defaults-from [[name] | none]

description [string]

activation-mode [none | allow | require]

edit smtps [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list smtps

list smtps [[[name] | [glob] | [regex]] ...]

show running-config smtps

show running-config smtps [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DELETE

delete smtps [name]

DESCRIPTION

You can use the smtps component to create, modify, display, or delete an SMTPs profile with which you can manage SMTPs traffic.

EXAMPLES

```
create smtps my_smtps_profile defaults-from smtps
```

Creates a custom SMTPs profile named my_smtps_profile that inherits its settings from the system default SMTPs profile.

```
list smtps
```

Displays the properties of all SMTPs profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is smtp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

activation-mode
Sets the activation-mode for STARTTLS. The options are NONE, ALLOW, or REQUIRE. The default value is NONE.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, sys provision, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013. All rights reserved.

BIG-IP 2013-11-19 ltm profile smtps(1)

Iitm profile socks

NAME

socks - Configures a SOCKS profile.

MODULE

ltm profile

SYNTAX

Configure the socks component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create socks [name]

modify socks [name]

options:

protocol-versions {
[[socks4] | [socks4a] | [socks5]] ... }

dns-resolver [dns-resolver]

ipv6 [no | yes]

tunnel-name [tunnel]

route-domain [route-domain]

default-connect-handling [deny | allow]

edit socks [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats socks

reset-stats socks [[[name] | [glob] | [regex]] ...]

DISPLAY

list socks

list socks [[[name] | [glob] | [regex]] ...]

show running-config socks

show running-config socks [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show socks

show socks [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete socks [name]

DESCRIPTION

You can use the socks component to create, modify, display, or delete an SOCKS profile.

The BIG-IP(r) system installation includes the following default SOCKS-type profiles:

socks

The default SOCKS profile contains values for properties related to managing SOCKS traffic.

You can create a new SOCKS-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

```
create socks my_socks_profile defaults-from socks
```

Creates a custom SOCKS profile named `my_socks_profile` that inherits its settings from the system default SOCKS profile.

OPTIONS

`protocols-versions`

Specifies the SOCKS protocol versions that are supported. The value is one or more off:

`socks4`

Specifies protocol support for SOCKS version 4.

`socks4a`

Specifies protocol support for SOCKS version 4A (like version 4, but with hostname support).

`socks5`

Specifies protocol support for SOCKS version 5 (with hostname and IPv6 support).

The default value specifies all available protocols.

`dns-resolver`

Specifies the `dns-resolver` object that will be used to resolve hostnames in connect requests. The default is `dns-resolver`.

`ipv6` Specifies the relative order of IPv4 and IPv6 DNS resolutions for URIs. The default is `no`, which will try a IPv4 lookup before a IPv6.

`tunnel-name`

Specifies the tunnel that will be used for outbound connect requests. This enables other virtual servers to receive connections initiated by the proxy service. The default is `socks-tunnel`.

`route-domain`

Specifies the `route-domain` that will be used for outbound connect requests. The default is `0`.

`default-connect-handling`

Specifies the behavior of the proxy service for connect requests. If set to `deny`, connect requests will only be honored if there is another virtual server listening for the requested outbound connection. If set to `allow` outbound connections will be made regardless of other virtual servers. The default is `deny`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `net dns-resolver`, `net route-domain`, `net tunnels`, `modify`, `regex`, `reset-stats`, `show`, `tmsk`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-07-11 ltm profile socks(1)

ltm profile splitsessionclient

NAME

`splitsessionclient` - Configures a `Splitsessionclient` profile.

MODULE

`splitsessionclient` profile

SYNTAX

Configure the `splitsessionclient` component within the `ltm profile` module using the syntax shown in the following sections.

CREATE/MODIFY

```
create splitsessionclient [name]
modify splitsessionclient [name]
options:
```

peer-ip
peer-port
http-header
local-peer
session-lookup-type

DISPLAY

list splitsessionclient
list splitsessionclient [[[name] | [glob] | [regex]] ...]
show running-config splitsessionclient
show running-config splitsessionclient [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties
 one-line
 partition

show splitsessionclient
show splitsessionclient [[[name] | [glob] | [regex]] ...]
options:
 (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
 field-fmt
 global

DELETE

delete splitsessionclient [name]

DESCRIPTION

Use this command to create, modify, display, or delete a Splitsessionclient profile.

EXAMPLES

create splitsessionclient my_splitsessionclient_profile defaults-from splitsessionclient

OPTIONS

peer-ip
Specifies the IP address of the Splitsession peer that the Out-of-band connection is made to if the splitsession peer (server) is remote. There is no default value for this field. If the peer is remote, splitsession-default-serverssl and splitsession-default-tcp profiles are used by the Splitsession client for the encrypted Out-of-band channel. If the splitsession peer is local, this value must be any.

peer-port
Specifies the port of the Splitsession peer that the Out-of-band connection is made to if the peer is remote. There is no default value for this field. If the peer is remote, splitsession-default-serverssl and splitsession-default-tcp profiles are used by the Splitsession client for the encrypted Out-of-band channel. If the splitsession peer is local, this value must be 0.

http-header
Specifies the HTTP header for signaling between the splitsession client and server. The default value is none.

local-peer
Specifies if the splitsession peer is local or remote. The default value is remote.

session-lookup-type
Specifies the session DB lookup method type, including flow, session-flow and http-header. Lookup type flow uses the 5-tuple flow key as the index to lookup the session DB, session-flow uses 5-tuple flow key plus the splitsessionclient profile name, and http-header depends on an HTTP header value. The default value is flow.

defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is none.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013-2016. All rights reserved.

Itm profile splitsessionserver

NAME

splitsessionserver - Configures a Splitsessionserver profile.

MODULE

splitsessionserver profile

SYNTAX

Configure the splitsessionserver component within the Itm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create splitsessionserver [name]

modify splitsessionserver [name]

options:

listen-ip

listen-port

http-header

local-peer

splitsessionclient

session-lookup-type

DISPLAY

list splitsessionserver

list splitsessionserver [[[name] | [glob] | [regex]] ...]

show running-config splitsessionserver

show running-config splitsessionserver [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show splitsessionserver

show splitsessionserver [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete splitsessionserver [name]

DESCRIPTION

Use this command to create, modify, display, or delete a Splitsessionserver profile.

EXAMPLES

create splitsessionserver my_splitsessionserver_profile defaults-from splitsessionserver

OPTIONS

listen-ip

Specifies the IP address that the Splitsession server listens on for the Out-of-band connection if the splitsession client is remote. There is no default value for this field. If the peer is remote, splitsession-default-clientssl and splitsession-default-tcp profiles are used by the Splitsession server for the encrypted Out-of-band channel. If the splitsession peer is local, this value must be any.

listen-port

Specifies the port that the Splitsession server listens on for the Out-of-band connection if the splitsession client is remote. There is no default value for this field. If the peer is remote, splitsession-default-clientssl and splitsession-default-tcp profiles are used by the Splitsession server for the encrypted Out-of-band channel. If the splitsession peer is local, this value must be 0.

http-header

Specifies the HTTP header for signaling between the splitsession client and server. The default value is none.

local-peer

Specifies if the splitsession peer is local or remote. The default value is remote.

splitsessionclient

Specifies splitsessionclient profile associated with this splitsessionserver in local deployment.

session-lookup-type

Specifies the session DB lookup method type, including flow, session-flow and http-header. Lookup type flow uses the 5-tuple flow key as the index to lookup the session DB, session-flow uses 5-tuple flow key plus the splitsessionclient profile name, and http-header depends on an HTTP header value. The default

value is flow.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is none.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013-2016. All rights reserved.

BIG-IP 2017-11-12 ltm profile splitsessionserver(1)

Itm profile statistics

NAME

statistics - Configures a Statistics profile.

MODULE

ltm profile

SYNTAX

Configure the statistics component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create statistics [name]

modify statistics [name]

options:

app-service [[string] | none]
defaults-from [[name] | none]
description [string]
field1 [string]
field2 [string]
field3 [string]
field4 [string]
field5 [string]
field6 [string]
field7 [string]
field8 [string]
field9 [string]
field10 [string]
field11 [string]
field12 [string]
field13 [string]
field14 [string]
field15 [string]
field16 [string]
field17 [string]
field18 [string]
field19 [string]
field20 [string]
field21 [string]
field22 [string]
field23 [string]
field24 [string]
field25 [string]
field26 [string]
field27 [string]
field28 [string]

field29 [string]
field30 [string]
field31 [string]
field32 [string]

edit statistics [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

mv statistics [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats statistics

reset-stats statistics [[[name] | [glob] | [regex]] ...]

DISPLAY

list statistics

list statistics [[[name] | [glob] | [regex]] ...]

show running-config statistics

show running-config statistics

[[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

show statistics

show statistics [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DELETE

delete statistics [all | name]

DESCRIPTION

You can use the statistics component to create, modify, display, or delete a Statistics profile that provides user-defined statistical counters.

EXAMPLES

```
create statistics my_stats_profile defaults-from stats
```

Creates a Statistics profile name `my_stats_profile` that inherits all settings and values from the profile `stats`.

```
list statistics my_stats
```

Displays the configuration of the profile `my_stats`.

```
list statistics my_stats field1 total_users field2 current_users field3 max_users
```

Creates a Statistics profile named `my_stats` with a total users counter in Field 1 and a current users counter in Field 2. You can then write an iRule to count the total number of connections, and record the current number of connections.

For more information about writing and using iRules(r), see the F5 Networks DevCentral web site at

.

```
mv statistics /Common/my_statistics_profile to-folder /Common/my_folder
```

Moves a custom statistics profile named `my_statistics_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is `stats`.

description

User defined description.

field1 ... field32

Specifies the name of a counter. You can specify a counter for up to 32 fields. The default value for each field is none.

You can then write an iRule that uses the counter names to gather statistics about the traffic the system is processing.

glob Displays the items that match the glob expression. See help `glob` for a description of glob expression

syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the component resides.

to-folder
statistics profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012. All rights reserved.

BIG-IP 2014-01-07 ltm profile statistics(1)

ltm profile stream

NAME

stream - Configures a Stream profile.

MODULE

ltm profile

SYNTAX

Configure the stream component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create stream [name]

modify stream [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

source [none | [string]]

target [none | [string]]

chunking-enabled [disabled | enabled]

chunk-size [integer]

edit stream [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv stream [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats stream

reset-stats stream [[[name] | [glob] | [regex]] ...]

DISPLAY

list stream

list stream [[[name] | [glob] | [regex]] ...]

show running-config stream

show running-config stream

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show stream

show stream [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete stream [name]

DESCRIPTION

You can use the stream component to search and replace strings within a data stream, such as a TCP connection.

EXAMPLES

```
create stream my_stream_profile defaults-from stream
```

Creates a custom Stream profile named `my_stream_profile` that inherits its settings from the system default stream profile.

```
list stream all-properties
```

Displays all properties for all Stream profiles.

```
mv stream /Common/my_stream_profile to-folder /Common/my_folder
```

Moves a custom stream profile named `my_stream_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is `stream`.

`description`

User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an `@` sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`source`

Specifies the string that you want to rewrite. The default value is none.

`target`

Specifies the new string, to replace the source string. The default value is none.

`chunking`

Specifies that the stream processor should parse incoming data in chunks. This restricts the maximum amount of additional memory needed by the stream processor. It also limits the maximum length of the value that can be matched to the `chunk-size`. The default value is disabled.

`chunk-size`

The maximum size of the parsed chunk.

`to-folder`

stream profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `mv`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2016-06-14 ltm profile stream(1)

Itm profile tcp-analytics

NAME

tcp-analytics - Configures a TCP Analytics profile.

MODULE

itm profile

SYNTAX

Configure the tcp-analytics component within the itm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create tcp-analytics [name]

modify tcp-analytics [name]

options:

collected-by-client-side [disabled | enabled]

collected-by-server-side [disabled | enabled]

collected-stats-external-logging [disabled | enabled]

collected-stats-internal-logging [disabled | enabled]

collect-city [disabled | enabled]

collect-continent [disabled | enabled]

collect-country [disabled | enabled]

collect-nexthop [disabled | enabled]

collect-post-code [disabled | enabled]

collect-region [disabled | enabled]

collect-remote-host-ip [disabled | enabled]

collect-remote-host-subnet [disabled | enabled]

defaults-from [[name] | none]

description [string]

external-logging-publisher [string]

edit tcp-analytics [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tcp-analytics

list tcp-analytics [[[name] | [glob] | [regex]] ...]

show running-config tcp-analytics

show running-config tcp-analytics

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show tcp-analytics

show tcp-analytics [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete tcp-analytics [name]

DESCRIPTION

You can use the tcp-analytics component to generate detailed TCP statistics for virtuals that have a TCP or FastL4 profile, if AVR is provisioned. The options specify where the data is collected and stored, and which connection identifiers ("entities") users can query the data with.

In general, storing more entities increases the ability to drill down to specific connections. However, if AVR hits its memory limits it will begin to aggregate new connections and reduce this resolution. If a virtual serves a very high number of IP addresses, for example, it is best to store data using collect-remote-host-subnet, or use the TCP::analytics iRule to enhance collection for a set of IP addresses of interest.

The system installation includes the default profile tcp-analytics. You can modify the settings of this profile, or create new TCP analytics profiles using it as a parent profile.

EXAMPLES

```
create itm profile tcp-analytics my_tcpanalytics_profile defaults-from tcp-analytics
```

Creates a custom TCP profile named my_tcpanalytics_profile that inherits its settings from the system default tcp-analytics profile.

```
list itm profile tcp-analytics all-properties
```

Displays all properties for all TCP profiles

OPTIONS

collected-by-client-side

When enabled, client side connections collect TCP analytics data unless directed otherwise by iRule. If false, client side connections only collect data if directed by iRule. The default value is enabled.

collected-by-server-side

When enabled, server side connections collect TCP analytics data unless directed otherwise by iRule. If false, server side connections only collect data if directed by iRule. The default value is disabled.

collected-stats-external-logging

When enabled, TCP statistics are logged on a remote machine determined by external-logging-publisher. The default value is disabled.

collected-stats-internal-logging

When enabled, TCP statistics are logged on the local BIG-IP. The default value is enabled.

collect-city

When enabled, AVR stores the city (from the GeoIP database) for the remote IP address of the connection with TCP statistics. See SOL11176 on support.f5.com for more on GeoIP. The default value is disabled.

collect-continent

When enabled, AVR stores the continent (from the GeoIP database) for the remote IP address of the connection with TCP statistics. See SOL11176 on support.f5.com for more on GeoIP. The default value is enabled.

collect-country

When enabled, AVR stores the country (from the GeoIP database) for the remote IP address of the connection with TCP statistics. See SOL11176 on support.f5.com for more on GeoIP. The default value is enabled.

collect-nexthop

When enabled, AVR stores the nexthop ethernet address of the connection with TCP statistics. The default value is disabled.

collect-post-code

When enabled, AVR stores the postcode (from the GeoIP database) for the remote IP address of the connection with TCP statistics. See SOL11176 on support.f5.com for more on GeoIP. The default value is disabled.

collect-region

When enabled, AVR stores the region (from the GeoIP database) for the remote IP address of the connection with TCP statistics. See SOL11176 on support.f5.com for more on GeoIP. The default value is disabled.

collect-remote-host-ip

When enabled, AVR stores the remote host IP address with TCP statistics. The default value is disabled.

collect-remote-host-subnet

When enabled, AVR stores the remote host IP subnet (24-bit) with TCP statistics. The default value is enabled.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is tcp-analytics.

description

User defined description.

external-logging-publisher

The name of the publisher to accept external logging.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-08-28 ltm profile tcp-analytics(1)

ltm profile tcp

NAME

tcp - Configures a Transmission Control Protocol (TCP) profile.

MODULE

ltm profile

SYNTAX

Configure the tcp component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create tcp [name]

modify tcp [name]

options:

abc [disabled | enabled]
ack-on-push [disabled | enabled]
app-service [[string] | none]
auto-proxy-buffer-size [disabled | enabled]
auto-receive-window-size [disabled | enabled]
auto-send-buffer-size [disabled | enabled]
close-wait-timeout [integer]
cmetrics-cache [disabled | enabled]
cmetrics-cache-timeout [integer]
congestion-control [high-speed | new-reno | none | reno | scalable |
vegas | illinois | woodside | chd | cdg | cubic | westwood | bbr]
defaults-from [[name] | none]
deferred-accept [disabled | enabled]
delay-window-control [disabled | enabled]
delayed-acks [disabled | enabled]
delay-window-control [disabled | enabled]
description [string]
dsack [disabled | enabled]
early-retransmit [disabled | enabled]
ecn [disabled | enabled]
enhanced-loss-recovery [disabled | enabled]
fast-open [disabled | enabled]
fast-open-cookie-expiration [integer]
fin-wait-timeout [integer]
fin-wait-2-timeout [integer]
hardware-syn-cookie [disabled | enabled]
idle-timeout [integer]
init-cwnd [integer]
init-rwnd [integer]
ip-tos-to-client [integer]
keep-alive-interval [integer]
limited-transmit [disabled | enabled]
link-qos-to-client [integer]
max-retrans [integer]
max-segment-size [integer]
md5-signature [disabled | enabled]
md5-signature-passphrase [none | [string]]
minimum-rto [integer]
mptcp [disabled | enabled | passthrough]
mptcp-csum [disabled | enabled]
mptcp-csum-verify [disabled | enabled]
mptcp-debug [disabled | enabled]
mptcp-fallback [reset | retransmit | active-accept | accept]
mptcp-join-max [integer]
mptcp-nojoindssack [disabled | enabled]
mptcp-rtomax [integer]
mptcp-rxmitmin [integer]
mptcp-subflowmax [integer]
mptcp-makeafterbreak [disabled | enabled]
mptcp-timeout [integer]
mptcp-fastjoin [disabled | enabled]
nagle [disabled | enabled | auto]
pkt-loss-ignore-rate [integer]
pkt-loss-ignore-burst [integer]
proxy-buffer-high [integer]
proxy-buffer-low [integer]
proxy-mss [disabled | enabled]
proxy-options [disabled | enabled]
push-flag [default | none | one | auto]
ip-df-mode [preserve | set | clear]
ip-ttl-mode [proxy | preserve | decrement | set]
ip-ttl-value [integer]
rate-pace [disabled | enabled]
rate-pace-max-rate [integer]
receive-window-size [integer]
reset-on-timeout [disabled | enabled]
rexmt-thresh [integer]
selective-acks [disabled | enabled]
selective-nack [disabled | enabled]
send-buffer-size [integer]
slow-start [disabled | enabled]
syn-cookie-enable [disabled | enabled]
syn-cookie-whitelist [disabled | enabled]
syn-max-retrans [integer]
syn-rto-base [integer]
tail-loss-probe [disabled | enabled]
time-wait-recycle [disabled | enabled]
time-wait-timeout [integer]
timestamps [disabled | enabled]
verified-accept [disabled | enabled]

zero-window-timeout [integer]

edit tcp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats tcp

reset-stats tcp [[[name] | [glob] | [regex]] ...]

DISPLAY

list tcp

list tcp [[[name] | [glob] | [regex]] ...]

show running-config tcp

show running-config tcp

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show tcp

show tcp [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete tcp [name]

DESCRIPTION

You can use the tcp component to manage TCP network traffic. Many of the options are standard SYSCTL-types of options, while others are unique to the traffic management system. For most of the options, the default values usually meet your needs. The specific options that you might want to change are: reset-on-timeout, idle-timeout, ip-tos-to-client, and link-qos-to-client.

The system installation includes these default TCP-type profiles: tc, tcp-cell-optimized, tcp-lan-optimized, and tcp-wan-optimized. You can modify the settings of these profiles, or create new TCP-type profiles using any of these existing profiles as parent profiles.

EXAMPLES

```
create tcp my_tcp_profile defaults-from tcp
```

Creates a custom TCP profile named my_tcp_profile that inherits its settings from the system default tcp profile.

```
list tcp all-properties
```

Displays all properties for all TCP profiles

OPTIONS

abc When enabled, increases the congestion window by basing the increase amount on the number of previously unacknowledged bytes that each acknowledgement code (ACK) includes. The default value is enabled.

ack-on-push

When enabled, significantly improves performance to Microsoft(r) Windows(r) and MacOS peers, who are writing out on a very small send buffer. The default value is enabled.

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

auto-proxy-buffer-size

Specifies, when enabled, that the system uses the network measurements to set the optimal proxy buffer size. The default value is disabled.

auto-receive-window-size

Specifies, when enabled, that the system uses the network measurements to set the optimal receive window size. The default value is disabled.

auto-send-buffer-size

Specifies, when enabled, that the system uses the network measurements to set the optimal send buffer size. The default value is disabled.

close-wait-timeout

Specifies the number of seconds that a connection remains in a LAST-ACK (last acknowledgement code) state before quitting. A value of 0 (zero) represents a term of forever (or until the maxrtx of the FIN state). The default value is 5 seconds.

cmetrics-cache

Specifies, when enabled, the default value, that the system uses a cache for storing congestion metrics.

cmetrics-cache-timeout

Specifies the time, in seconds, for which entries in the congestion metrics cache are valid. The default

value is 0, which defers to the sys db variable route.metrics.timeout.

congestion-control

Specifies the algorithm to use to share network resources among competing users to reduce congestion. The default value is high-speed.

The options are:

bbr Specifies that the system uses an implementation of the BBR congestion control algorithm, which uses connection's maximum delivery rate (bottleneck bandwidth) and minimum round trip time to avoid growing buffers and path delay.

cdg Specifies that the system use a Caia Delay-Gradient congestion control algorithm, where congestion inferences are made based on a gradient of RTT over time. Improves inferences made about packet loss and whether they are due to congestion or other factors. The use of a shadow window improves coexistence with loss-based TCP flows.

chd Specifies that the system use a Caia-Hamilton delay-based congestion control algorithm, where delay-based congestion window operations are performed only once per RTT. Tolerates packet losses that are likely to be unrelated to congestion. Uses a shadow window to help regain lost transmission opportunities when competing with loss-based TCP flows.

cubic Specifies that the system uses a component optimized for high latency, high bandwidth connections as the TCP congestion control algorithm.

high-speed

Specifies that the system uses a more aggressive, loss-based algorithm.

illinois

Specifies that the system uses a hybrid of both delay and loss as the TCP congestion control algorithm.

new-reno

Specifies that the system uses a modification to the Reno algorithm that responds to partial acknowledgements when SACKs are unavailable.

none

Specifies that the system does not use a network-congestion-control mechanism, even when congestion occurs.

reno

Specifies that the system uses an implementation of the TCP Fast Recovery algorithm, which is based on the implementation in the BSD Reno release.

scalable

Specifies that the system uses a TCP algorithm modification that adds a scalable, delay-based and loss-based component into the Reno algorithm.

vegas

Specifies that the system uses a delay-based component as the TCP congestion control algorithm.

westwood

Specifies that the system uses the Westwood+ bandwidth estimation component as the TCP congestion control algorithm.

woodside

Specifies that the system uses a hybrid of both delay and loss as the TCP congestion control algorithm.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is tcp.

deferred-accept

Specifies, when enabled, that the system defers allocation of the connection chain context until the system has received the payload from the client. This option is useful for dealing with 3-way handshake denial-of-service (DOS) attacks. The default value is disabled.

delay-window-control

When enabled, the system uses an estimate of queueing delay as a measure of congestion, in addition to the normal loss-based control, to control the amount of data sent. The default value is disabled.

delayed-acks

Specifies, when enabled, the default value, that the traffic management system allows coalescing of multiple acknowledgement (ACK) responses.

description

User defined description.

dsack

When enabled, specifies the use of the SACK option to acknowledge duplicate segments. The default is disabled.

early-retransmit

Specifies, when enabled, that the system uses early retransmit recovery (as specified in RFC 5827) to reduce the recovery time for connections that are receive-buffer or user-data limited. The default value is enabled.

ecn Specifies, when enabled, that the system uses the TCP flags CWR and ECE to notify its peer of congestion and congestion counter-measures. The default value is enabled.

enhanced-loss-recovery

Specifies whether the system uses enhanced loss recovery to recover from random packet losses more effectively. The default value is enabled.

fast-open

Specifies, when enabled, that the system supports TCP Fast Open, which allows a client to include the first packet of data with the SYN to reduce latency. The default value is enabled. This option has no effect on server-side TCP profiles.

fast-open-cookie-expiration

Specifies the number of seconds that a "Fast Open Cookie" delivered to a client is valid for SYN packets from that client. The default value is 21600 seconds (6 hours). A value of 0 (zero) means use the default. The maximum value is 1000000 seconds.

fin-wait-timeout

Specifies the number of seconds that a connection is in the FIN-WAIT-1 or closing state before quitting. The default value is 5 seconds. A value of 0 (zero) represents a term of forever (or until the maxrtx of the FIN state).

fin-wait-2-timeout

Specifies the number of seconds that a connection is in the FIN-WAIT-2 state before quitting. The default value is 300 seconds. A value of 0 (zero) represents a term of forever (or until the maxrtx of the FIN state).

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

hardware-syn-cookie

This option is deprecated in version 13.0.0 and is replaced by syn-cookie-enable. Specifies whether or not to use hardware SYN Cookie when cross system limit. The default value is enabled.

idle-timeout

Specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 300 seconds.

init-cwnd

Specifies the initial congestion window size for connections to this destination. The actual window size is this value multiplied by the MSS (Maximum Segment Size) for the same connection. The default value is 10. The range is from 0 to 64.

init-rwnd

Specifies the initial receive window size for connections to this destination. The actual window size is this value multiplied by the MSS (Maximum Segment Size) for the same connection. The default value is 10. The range is from 0 to 64.

ip-df-mode

Describe the IP Header Don't Fragment (DF) bit setting in the outgoing TCP packet. The available settings are: Pmtu: Set the outgoing IP Header DF bit based on IP pmtu setting(tm.pathmtudiscovery). Preserve: Set the outgoing Packet's IP Header DF bit to be same as incoming IP Header DF bit. Set: Set the outgoing packet's IP Header DF bit. Clear: Clear the outgoing packet's IP Header DF bit. The default setting is Pmtu.

ip-ttl-mode

Describe the outgoing TCP packet's IP Header TTL mode. The available Modes are: Proxy: Set the outgoing IP Header TTL value to 255/64 for ipv4/ipv6 respectively. Preserve: Set the outgoing IP Header TTL value to be same as the incoming IP Header TTL value. Decrement: Set the outgoing IP Header TTL value to be one less than the incoming TTL value. Set: Set the outgoing IP Header TTL value to a specific value(as specified by ip-ttl-v[4|6]). The default mode is Proxy.

ip-ttl-v4

Specify the outgoing packet's IP Header TTL value for IPv4 traffic. Maximum TTL value that can be specified is 255. The default is 255.

ip-ttl-v6

Specify the outgoing packet's IP Header TTL value for IPv6 traffic. Maximum TTL value that can be specified is 255. The default is 64.

ip-tos-to-client

Specifies the Type of Service (ToS) level that the traffic management system assigns to TCP packets when sending them to clients. The default value is 0 (zero).

keep-alive-interval

Specifies the keep-alive probe interval, in seconds. The default value is 1800 seconds.

limited-transmit

Specifies, when enabled, the default value, that the system uses limited transmit recovery revisions for fast retransmits (as specified in RFC 3042) to reduce the recovery time for connections on a lossy network.

link-qos-to-client

Specifies the Link Quality of Service (QoS) level that the system assigns to TCP packets when sending them to clients. The default value is 0 (zero).

max-retrans

Specifies the maximum number of retransmissions of data segments that the system allows. The default value is 8.

max-segment-size

Specifies the largest amount of data that the system can receive in a single TCP segment, not including the TCP and IP headers. If the value is 0 (zero), the system calculates the value from the MTU. The default value is 1460 bytes.

md5-signature

Specifies, when enabled, that the system uses RFC2385 TCP-MD5 signatures to protect TCP traffic against intermediate tampering. The default value is disabled.

md5-signature-passphrase

Specifies a plain text passphrase which may be between 1 and 80 characters in length, and is used in a shared-secret scheme to implement the spoof-prevention parts of RFC2385. The default value is none.

minimum-rto

Specifies the minimum TCP retransmission timeout in milliseconds. The default value is 1000 milliseconds.

mptcp

Specifies, when enabled, that the system will accept MPTCP connections. When passthrough MPTCP connections are not terminated by this virtual. The default value is disabled.

mptcp-csum

Specifies, when enabled, that the system will calculate the checksum for MPTCP connections. The default value is disabled.

mptcp-csum-verify

Specifies, when enabled, that the system verifies checksum for MPTCP connections. The default value is disabled.

mptcp-debug

This option is DEPRECATED v12.0.0 onwards and is maintained here for backward compatibility reasons. Specifies, when enabled, that the system provides debug logs and statistics for MPTCP connections. The default value is disabled.

mptcp-fallback

Specifies, MPTCP fallback mode. The default value is reset.

The options are:

accept

Specifies accept on fallback.

active-accept

Specifies active accept on fallback.

reset

Specifies that the connection is reset on fallback.

retransmit

Specifies retransmit on fallback.

mptcp-join-max

Specifies the max number of MPTCP connections that can join to given one. The default value is 5.

mptcp-nojoindssack

Specifies, when enabled, no DSS option is sent on the JOIN ACK. The default value is disabled.

mptcp-rtomax

Specifies, the number of RTOs before declaring subflow dead. The default value is 5.

mptcp-rxmitmin

Specifies the minimum value (in msec) of the retransmission timer for these MPTCP flows. The default value is 1000.

mptcp-subflowmax

Specifies the maximum number of MPTCP subflows for a single flow. The default value is 6.

mptcp-makeafterbreak

Specifies, when enabled, that make-after-break functionality is supported, allowing for long-lived MPTCP sessions. The default value is disabled.

mptcp-timeout

Specifies, the timeout value to discard long-lived sessions that do not have an active flow, in seconds. The default value is 3600.

mptcp-fastjoin

Specifies, when enabled, FAST join, allowing data to be sent on the MP_JOIN SYN, which can allow a server response to occur in parallel with the JOIN. The default value is disabled.

nagle

Specifies, when enabled, that the system applies Nagle's algorithm to reduce the number of short segments on the network. The default value is disabled. When auto, the use of Nagle's algorithm is decided based on network conditions.

Note that for interactive protocols such as Telnet, rlogin, or SSH, F5 Networks recommends disabling this setting on high-latency networks, to improve application responsiveness.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition
Displays the administrative partition within which the profile resides.

pkt-loss-ignore-burst
Specifies the probability of performing congestion control when multiple packets in a row are lost, even if the pkt-loss-ignore-rate was not exceeded. Valid values are 0 (zero) through 32. The default value is 0 (zero), which means that the system performs congestion control, if any packets are lost. Higher values decrease the chance of performing congestion control.

pkt-loss-ignore-rate
Specifies the threshold of packets lost per million at which the system should perform congestion control. Valid values are 0 (zero) through 1,000,000. The default value is 0 (zero), which means that the system performs congestion control, if any packet loss occurs. If you set the ignore rate to 10 and packet loss for a TCP connection is greater than 10 per million, congestion control occurs.

proxy-buffer-high
Specifies the highest level at which the receive window is closed. The default value is 131072.

proxy-buffer-low
Specifies the lowest level at which the receive window is closed. The default value is 98304.

proxy-mss
Specifies, when enabled, that the system advertises the same TCP maximum segment size to the server as was negotiated with the client. The setting is ignored when MRF routing (e.g., httprouter, siprouter, diameterrouter, mqttrouter, messagerouter) is used. The default value is enabled.

proxy-options
Specifies, when enabled, that the system advertises an option, such as a time-stamp to the server only if it was negotiated with the client. The setting is ignored when MRF routing (e.g., httprouter, siprouter, diameterrouter, mqttrouter, messagerouter) is used. The default value is disabled.

push-flag
When default, specifies that the system sets PUSH flag when sending the last segment in the send buffer. When none, specifies that the system never sets PUSH flag for TCP packets. When one, specifies that the system sets one PUSH flag for the FIN segment. When auto, specifies that the system sets PUSH flag based on the application/network conditions. The default value is default.

rate-pace
Specifies, when enabled, that the system will rate pace TCP data transmissions. The default value is enabled.

rate-pace-max-rate
If not 0, sets the maximum rate in bytes per second that TCP data transmission will be paced to. If set to 0, no maximum is enforced. The default value is 0.

receive-window-size
Specifies the size of the receive window, in bytes. The default value is 65535 bytes.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reset-on-timeout
Specifies whether to reset connections on timeout. The default value is enabled.

rexmt-thresh
Specifies the number of duplicate ACKs (retransmit threshold) to start fast recovery. The default value is 3. The range is from 3 to 255.

selective-acks
Specifies, when enabled, the default value, that the system negotiates RFC2018-compliant Selective Acknowledgements with peers.

selective-nack
Specifies whether Selective Negative Acknowledgment is enabled or disabled. The default value is disabled.

send-buffer-size
Specifies the size of the buffer, in bytes. The default value is 131072 bytes.

slow-start
Specifies, when enabled, the default value, that the system uses larger initial window sizes (as specified in RFC 3390) to help reduce round trip times. Note that disabling this attribute causes the setting for cmetrics-cache to be ignored.

syn-cookie-enable
Specifies the default (if no DoS profile is associated) number of embryonic connections that are allowed on any virtual server, before SYN Cookie challenges are enabled for that virtual server. The default is enabled.

syn-cookie-whitelist
Specifies whether or not to use a SYN Cookie WhiteList when doing software SYN Cookies. This means not doing a SYN Cookie for the same src IP address if it has been done already in the previous

tm.flowstate.timeout (30) seconds. The default value is disabled.

syn-max-retrans

Specifies the maximum number of retransmissions of SYN segments that the system allows. The default value is 3.

syn-rto-base

Specifies the initial RTO (Retransmission TimeOut) base multiplier for SYN retransmission, in milliseconds. This value is modified by the exponential backoff table to select the interval for subsequent retransmissions. The default value is 3000.

tail-loss-probe

Specifies whether the system uses tail loss probe to reduce the number of retransmission timeouts. The default value is enabled.

tcp-options

Specifies the option numbers that will be accessible from iRules (TCP::option) for the flow. The format of each entry should be: "{

ltm profile tftp

NAME

tftp - Configures a TFTP profile.

MODULE

ltm profile

SYNTAX

Configure the tftp component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create tftp [name]

modify tftp [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

idle-timeout [integer]

log-publisher [log publisher name | none]

log-profile [log profile name | none]

mv tftp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

edit tftp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list tftp

list tftp [[[name] | [glob] | [regex]] ...]

show running-config tftp

show running-config tftp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete tftp [name]

DESCRIPTION

Use this command to create, modify, display, or delete an TFTP profile with which you can manage TFTP traffic.

EXAMPLES

create tftp my_tftp_profile defaults-from tftp

Creates a custom TFTP profile named my_tftp_profile that inherits its settings from the system default TFTP profile.

list tftp

Displays the properties of all TFTP profiles.

mv tftp /Common/my_tftp_profile to-folder /Common/my_folder

Moves a custom tftp profile named `my_tftp_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is `tftp`.

`description`

User defined description.

`glob`

Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`to-folder`

tftp profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

`idle-timeout`

Specifies an idle timeout in seconds. This setting specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 300 seconds.

`log-publisher`

Specify the name of the log publisher which logs translation events. See help `sys log-config` for more details on the logging sub-system. Use the "sys log-config publisher" component to set up a log publisher.

`log-profile`

Specify the name of the ALG log profile which controls the logging of ALG events. See help `itm alg-log-profile` for more details on the logging profile sub-system. Use the "itm alg-log-profile profile" component to set up an ALG log profile.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `itm virtual`, `modify`, `mv`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 itm profile tftp(1)

Itm profile traffic-acceleration

NAME

`traffic-acceleration` - Configures a Traffic Acceleration profile.

MODULE

itm profile

SYNTAX

Configure the `traffic-acceleration` component within the `itm profile` module using the syntax shown in the following sections.

CREATE/MODIFY

`create traffic-acceleration [name]`

`modify traffic-acceleration [name]`

options:

`app-service` [[string] | none]

`defaults-from` [[name] | none]

`description` [string]

```
idle-timeout [integer]
tcp-handshake-timeout [integer]
time-wait-timeout [integer]
```

```
edit traffic-acceleration [ [name] | [regex] ]
options:
  all-properties
  non-default-properties
```

```
reset-stats traffic-acceleration
reset-stats traffic-acceleration [ [name] | [regex] ]
```

```
DISPLAY
list traffic-acceleration
list traffic-acceleration [ [name] | [regex] ]
show running-config traffic-acceleration
show running-config traffic-acceleration [ [name] | [regex] ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show traffic-acceleration
show traffic-acceleration [ [name] | [regex] ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global
```

```
DELETE
delete traffic-acceleration [name]
```

DESCRIPTION

You can use this component to create, modify, display, or delete a Traffic Acceleration profile.

Any changes you make to an active Traffic Acceleration profile (one that is in use by a virtual server) take effect after the value of the idle-timeout option has passed. That means new connections are affected by the profile change immediately. However, for the new values to take effect, old connections need to be either aged out or closed.

EXAMPLES

```
create traffic-acceleration my_profile defaults-from traffic-acceleration
```

Creates a custom Traffic Acceleration profile named my_profile that inherits its settings from the system default Traffic Acceleration profile.

```
show traffic-acceleration
```

Displays statistics for all Traffic Acceleration profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is traffic-acceleration.

description

User defined description.

idle-timeout

Specifies the number of milliseconds that a connection is idle before the connection is eligible for deletion. The default value is five minutes or 300000 milliseconds. Minimum is 0 meaning the connection does not timeout. The maximum is 2147483647 (0x7FFFFFFF) milliseconds.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

tcp-handshake-timeout

Specifies a TCP handshake timeout in milliseconds. The default value is 5000 milliseconds. Minimum is 0 meaning there is no timeout. The maximum is 65535 milliseconds.

time-wait-timeout

Specifies a TCP time_wait timeout in milliseconds. The default value is 2000 milliseconds. The minimum is 1 millisecond. The maximum is 65535 milliseconds.

SEE ALSO

create, delete, edit, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2016-10-31 ltm profile traffic-acceleration(1)

ltm profile udp

NAME

udp - Configures a User Datagram Protocol (UDP) profile.

MODULE

ltm profile

SYNTAX

Configure the udp component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create udp [name]

modify udp [name]

options:

allow-no-payload [disabled | enabled]

app-service [[string] | none]

buffer-max-bytes [integer]

buffer-max-packets [integer]

datagram-load-balancing [disabled | enabled]

defaults-from [[name] | none]

description [string]

idle-timeout [immediate | indefinite | integer]

ip-tos-to-client [[integer] | pass-through]

link-qos-to-client [[integer] | pass-through]

no-checksum [disabled | enabled]

proxy-mss [disabled | enabled]

ip-ttl-mode [proxy | preserve | decrement | set]

ip-ttl-v4 [integer]

ip-ttl-v6 [integer]

ip-df-mode [pmtu | preserve | set | clear]

send-buffer-size [integer]

edit udp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv udp [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

reset-stats udp

reset-stats udp [[[name] | [glob] | [regex]] ...]

DISPLAY

list udp

list udp [[[name] | [glob] | [regex]] ...]

show running-config udp

show running-config udp

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

show udp

show udp [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete udp [name]

DESCRIPTION

You can use the udp component to manage UDP network traffic.

EXAMPLES

create udp my_udp_profile defaults-from udp

Creates a custom UDP profile named `my_udp_profile` that inherits its settings from the system default UDP profile.

`list udp all-properties`

Displays all properties for all UDP profiles.

`mv udp /Common/my_udp_profile to-folder /Common/my_folder`

Moves a custom udp profile named `my_udp_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`allow-no-payload`

Provides the ability to allow the passage of datagrams that contain header information, but no essential data. The default is disabled.

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`buffer-max-bytes`

Specifies ingress buffer byte limit. The default value is 655350. Maximum allowed value is 16777215.

`buffer-max-packets`

Specifies ingress buffer packet limit. The default value is 0. Maximum allowed value is 255.

`datagram-load-balancing`

Provides the ability to load balance UDP datagram by datagram. The default is disabled.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is `udp`.

`description`

User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`idle-timeout`

Specifies the number of seconds that a connection is idle before the connection is eligible for deletion. The default value is 60 seconds.

`ip-tos-to-client`

Specifies the Type of Service level that the traffic management system assigns to UDP packets when sending them to clients. The default value is 0 (zero).

`ip-ttl-mode`

Describe the outgoing UDP packet's TTL mode. Modes are: Proxy: Set the IP TTL of ipv4 to the default value of 255 and ipv6 to the default value of 64. Preserve: Set the IP TTL to the original packet TTL value. Decrement: Set the IP TTL to the original packet TTL value minus 1. Set: Set the IP TTL with the specified values in `ip-ttl-v4` and `ip-ttl-v6` values in the same profile.

`ip-df-mode`

Describe the Don't Fragment (DF) bit setting in the outgoing UDP packet. Pmtu: Set the outgoing UDP packet DF big based on the ip pmtu setting. Preserve: Preserve the incoming UDP packet Don't Fragment bit. Set: Set the outgoing UDP packet DF bit. Clear: Clear the outgoing UDP packet DF bit.

`link-qos-to-client`

Specifies the Quality of Service level that the system assigns to UDP packets when sending them to clients. The default value is 0 (zero).

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`no-checksum`

Enables or disables checksum processing. Note that if the datagram is IPv6, the system always performs checksum processing. The default value is disabled.

`partition`

Displays the administrative partition within which the profile resides.

`proxy-mss`

Specifies, when enabled, that the system advertises the same mss to the server as was negotiated with the client. The default value is disabled.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`send-buffer-size`

Specifies send buffer byte limit. The default value is 655350. The range is from 536 to 16777215.

to-folder
udp profiles can be moved to any folder under /Common, but configuration dependencies may restrict moving the profile out of /Common.

SEE ALSO

create, delete, edit, glob, ltm profile, ltm virtual, modify, mv, show, regex, reset-stats, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2017-07-24 ltm profile udp(1)

ltm profile wa-cache

NAME

wa-cache - Manages the BIG-IP(r) system WebAccelerator cache.

MODULE

ltm profile

SYNTAX

Configure the wa-cache component within the ltm profile module using the syntax shown in the following sections.

DELETE

delete wa-cache [name]

DESCRIPTION

You can use the wa-cache component to delete the entries in the BIG-IP(r) system WebAccelerator cache.

EXAMPLES

delete wa-cache

Deletes the entries in the BIG-IP system WebAccelerator cache.

SEE ALSO

delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-10-19 ltm profile wa-cache(1)

ltm profile web-acceleration

NAME

web-acceleration - Configures a Web Acceleration profile.

MODULE

ltm profile

SYNTAX

Configure the web-acceleration component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create web-acceleration [name]

modify web-acceleration [name]

options:

app-service [[string] | none]

cache-aging-rate [integer]

cache-client-cache-control-mode [all | max-age | none]

cache-insert-age-header [disabled | enabled]

cache-max-age [integer]

```

cache-max-entries [integer]
cache-object-max-size [integer]
cache-object-min-size [integer]
cache-size [integer]
cache-uri-exclude
  [add | delete | replace-all-with] {
[URI] ...
  }
cache-uri-exclude none
cache-uri-include
  [add | delete | replace-all-with]{
[URI] ...
  }
cache-uri-include .*
cache-uri-include-override
  [add | delete | replace-all-with]{
[URI] ...
  }
cache-uri-include-override none
cache-uri-pinned
  [add | delete | replace-all-with] {
[URI] ...
  }
cache-uri-pinned none
defaults-from [ [name] | none]
description [string]

```

```

edit web-acceleration [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

```

reset-stats web-acceleration
reset-stats web-acceleration [ [ [name] | [glob] | [regex] ] ... ]

```

```

DISPLAY
list web-acceleration
list web-acceleration [ [ [name] | [glob] | [regex] ] ... ]
show running-config web-acceleration
show running-config web-acceleration [ [ [name] | [glob] | [regex] ]
... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

```

```

show web-acceleration
show web-acceleration [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global

```

```

DELETE
delete web-acceleration [name]

```

DESCRIPTION

You can use the web-acceleration component to create, modify, display, or delete an Web Acceleration profile.

The BIG-IP(r) system installation includes the following default Web Acceleration-type profiles:

```

web-acceleration
optimized-caching
optimized-acceleration

```

The default Web Acceleration profile contains values for properties related to managing WA Cache.

You can create a new Web Acceleration-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

```

create web-acceleration my_wa_profile defaults-from web-acceleration

```

Creates a custom Web Acceleration profile named my_wa_profile that inherits its settings from the system default Web Acceleration profile.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

cache-aging-rate

Specifies how quickly the system ages a cache entry. The aging rate ranges from 0 (slowest aging) to 10 (fastest aging). The default value is 9.

cache-client-cache-control-mode

Specifies which cache disabling headers sent by clients the system ignores. The default value is all.

cache-insert-age-header

When enabled, inserts Age and Date headers in the response. The default value is enabled.

cache-max-age

Specifies how long the system considers the cached content to be valid. The default value is 3600 seconds.

cache-max-entries

Specifies the maximum number of entries that can be in the WA cache. The default value is 10000.

cache-object-max-size

Specifies the largest object that the system considers eligible for caching. The default value is 50000 bytes.

cache-object-min-size

Specifies the smallest object that the system considers eligible for caching. The default value is 500 bytes.

cache-size

Specifies the maximum size, in megabytes, for the WA cache. When the cache reaches the maximum size, the system starts removing the oldest entries. The default value is 100 megabytes.

cache-uri-exclude

Configures a list of Uniform Resource Identifiers (URIs) to exclude from the WA Cache. The default value is none and specifies that no URI will be excluded.

cache-uri-include

Configures a list of URIs that are cacheable. The default value is .* and specifies that all URIs are cacheable.

cache-uri-include-override

Configures a list of URIs that should be cached in the WA cache even though they would normally not be cached due to constraints defined by cache-object-max-size or others. The default value is none. URIs on the cache-uri-include-override list are cacheable even if they are not on the cache-uri-include list.

cache-uri-pinned

Configures a list of URIs that are kept in the WA cache regardless their max-age or expiry settings. The default value is none. URIs on the cache-uri-pinned list are cacheable even if they are not on the cache-uri-include list.

defaults-from

Configures the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is webacceleration.

description

User defined description.

partition

Displays the administrative partition within which the profile resides.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2020-01-07 ltm profile web-acceleration(1)

ltm profile web-security

NAME

web-security - Configures a Web Security profile.

MODULE

ltm profile

SYNTAX

Configure the web-security component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create web-security [name]

modify web-security [name]
options:
defaults-from [[name] | none]

edit web-security [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list web-security
list web-security [[[name] | [glob] | [regex]] ...]
...]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete web-security [name]

DESCRIPTION

You can use the web-security component to create, modify, display, or delete an Web Security profile.

The BIG-IP(r) system installation includes the following default Web Security-type profiles:

websecurity

The default Web Security profile contains values for properties related to managing web security.

You can create a new Web Security-type profile using an existing profile as a parent profile, and then you can change the values of the properties to suit your needs.

EXAMPLES

create web-security my_asm_profile defaults-from web-security

Creates a custom Web Security profile named my_asm_profile that inherits its settings from the system default Web Security profile.

OPTIONS

defaults-from

Configures the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is none.

partition

Displays the administrative partition within which the profile resides.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 ltm profile web-security(1)

Itm profile websocket

NAME
websocket - Configures a WebSocket protocol profile.

MODULE
websocket profile

SYNTAX
Configure the websocket component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY
create websocket [name]
modify websocket [name]
options:
masking [preserve | unmask | remark | selective]
compress-mode [preserved | typed]
compression [enabled | disabled]

window-bits [integer]
no-delay [enabled | disabled]

edit websocket [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

mv websocket [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-
folder [folder-name]]]
options:
to-folder

reset-stats websocket
reset-stats websocket [[[name] | [glob] | [regex]] ...]

DISPLAY
list websocket
list websocket [[[name] | [glob] | [regex]] ...]
show running-config websocket
show running-config websocket [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

show websocket
show websocket [[[name] | [glob] | [regex]] ...]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
global

DELETE
delete websocket [name]

DESCRIPTION
Use this command to create, modify, display, or delete a WebSocket profile.

EXAMPLES
create websocket my_websocket_profile defaults-from websocket

OPTIONS
masking
Specifies the masking operation for WebSocket frames.

preserve
Maintains the mask of the packet received from the client.

unmask
Removes the mask from the packet and remasks it using the same mask when sending the traffic to the server.

remask
Removes the mask received from the client and the system generates a new, random mask when sending the traffic to the server.

selective
Preserves the mask of the packet received, and makes no changes unless an Application Security Policy is associated with the virtual server.

compress-mode
Specifies the mode that controls what compression operations are performed. The default value is preserved.

preserved
No compression or decompression operations are performed in this mode.

typed
The compression parameter controls whether compression is negotiated with the endpoint. Compressed data received from the endpoint is decompressed thereby enabling security policies to take actions on cleartext data. Data maybe recompressed before being masked and sent on the wire.

= over 5

compression
Specifies whether compression will be negotiated with the endpoint. The default value is enabled.

window-bits
Specifies the maximum sliding window for compression negotiated with the endpoint. The default value is 10.

no-delay
Specifies whether data should be buffered for efficient compression, or compressed without delay. The default value is enabled.

= back

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is none.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm profile fasthttp, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013-2016. All rights reserved.

BIG-IP 2019-03-29 ltm profile websocket(1)

Itm profile xml

NAME

xml - Configures an XML profile.

MODULE

ltm profile

SYNTAX

Configure the xml component within the ltm profile module using the syntax shown in the following sections.

CREATE/MODIFY

create xml [name]

modify xml [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

namespace-mappings [[none] |

[{ { mapping-namespace namespace1 mapping-prefix prefix1 } }]

xpath-queries [none |

[add | delete | replace_all_with { queries }]]

multiple-query-matches [enabled | disabled]

edit xml [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

mv xml [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list xml

list xml [[[name] | [glob] | [regex]] ...]

show running-config xml

show running-config xml [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete xml [name]

DESCRIPTION

Use this command to create, modify, display, or delete an XML profile with which you can use XML functionality.

EXAMPLES

```
create xml my_xml_profile defaults-from xml
```

Creates a custom XML profile named `my_xml_profile` that inherits its settings from the system default XML profile.

```
list xml
```

Displays the properties of all XML profiles.

```
mv xml /Common/my_xml_profile to-folder /Common/my_folder
```

Moves a custom xml profile named `my_xml_profile` to a folder named `my_folder`, where `my_folder` has already been created and exists within `/Common`.

OPTIONS

`app-service`

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

`defaults-from`

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is `xml`.

`description`

User defined description.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`

Displays the administrative partition within which the component resides.

`namespace-mappings`

Specifies a list of mappings between namespaces and prefixes to be used in the XPath queries of the profile.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`to-folder`

xml profiles can be moved to any folder under `/Common`, but configuration dependencies may restrict moving the profile out of `/Common`.

`xpath-queries`

Specifies the list of XPath queries that are used by the profile. A match of any of the queries will trigger the `XML_CONTENT_BASED_ROUTING` iRule event.

`multiple-query-matches`

Enables or disables multiple matches for a single XPath query.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `mv`, `regex`, `reset-stats`, `show`, `tms`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013. All rights reserved.

BIG-IP 2014-01-07 ltm profile xml(1)

ltm rule-profiler

MODULE

ltm

SYNTAX

Configure the rule-profiler component within the ltm module using the syntax shown in the following sections.

```
CREATE/MODIFY
create rule-profiler [name]
modify rule-profiler [ [name] ... ]
```

```
DISPLAY
list rule-profiler
list rule-profiler [ [ [name] | [glob] | [regex] ] ... ]
```

```
DELETE
delete rule-profiler [name]
```

```
HELP help rule-profiler
```

DESCRIPTION

You can use the rule-profiler component to configure execution tracing of iRules attached to virtual servers. A rule-profiler comprises a set of matching filters that help limit the profiling information output, a log publisher that directs the output to the intended destination and properties for controlling the state of the profiler and the amount of time the profiler will be active for.

EXAMPLES

```
list rule-profiler
```

Displays all iRule Profiler objects.

```
delete rule-profiler my_profiler
```

Deletes the iRule Profiler named my_profiler.

```
rule-profiler my_profiler {
vs-filter { vs_test }
rule-filter { r1_http }
event-filter { HTTP_REQUEST }
```

```
period 1000
publisher my_tracer_pub
occ-mask 30
state disabled
}
```

Creates an iRule Profiler named my_profiler.

```
modify rule-profiler my_profiler state enabled
```

Enables the rule-profiler causing matching profiling occurrences to be published by the publisher associated with the rule-profiler.

OPTIONS

description
User defined description.

period
Specifies the period (in milliseconds) that a rule-profiler can be active for. When set to 0, the profiler is disabled. Default value is 0.

state
Specifies the state of the rule-profiler. When set to enabled and the profiler consequently started, matching occurrence records are emitted via the configured log publisher. Default value is disabled.

start
Start profiling. After this command, at the point of the first filter match, profiling and timer counting start.

stop Stop profiling. After this command, the profiling stops.

vs-filter
Specifies the name of virtual server(s) for which occurrence records can be emitted. If this property is left empty, all virtual servers will be considered a match. Default is empty list.

rule-filter
Specifies the name of iRule(s) for which occurrence records can be emitted. If this property is left empty, all executing iRules will be considered a match. Default is empty list.

event-filter
Specifies the name of iRule Events for which occurrence records can be emitted. If this property is left empty, all iRule Events will be considered a match. Default is empty list.

occ-mask
Specifies the types of occurrences of interest to the user. This is a value set comprised by adding the desired occurrence type. The types are defined as follows:

• event

Specifies that profiling information is emitted at the entry point to the TCL subsystem, as close as possible to the occurrence of the external event, and at completion of the execution, just before the TCL subsystem returns to the LTM profile/filter code.

• rule

Specifies that profiling information is emitted at the entry point to the execution of a single iRule handler for the event, and once the TCL VM completes execution of the single event handler. This bit is useful when there are multiple rules handling the same event.

• rule-vm

Specifies that profiling information is emitted once the TCL VM is about to execute the byte code for an iRule script, and once it completed executing the byte code.

• cmd-vm

Specifies that profiling information is emitted when the TCL VM is beginning the internal setup preceding the invocation of a TCL command (eg. HTTP::uri) and once the internal setup is being torn down.

• cmd

Specifies that profiling information is emitted once the TCL VM is about to invoke a TCL command (eg. HTTP::uri) and once the command execution completes.

• var-mod

Specifies that profiling information is emitted when the TCL VM is changing the value of an iRule variable.

• bytecode

Specifies that profiling information is emitted when the TCL VM is executing a byte code instruction.

SEE ALSO

create, delete, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-05-12 Itm rule-profiler(1)

Itm rule

NAME

rule - Configures an iRule for traffic management system configuration.

MODULE

itm

SYNTAX

Configure the rule component within the Itm module using the syntax shown in the following sections.

CREATE/MODIFY

create rule [name]

edit rule [name]

modify rule [[name] | [glob] | [regex]] ...]

Note: When using tmsh, you can only create iRules using the editor, which starts when you use the create or edit commands. You cannot create an iRule directly on the command line. The vim editor applies the autoindent and smartindent options. You can toggle on/off paste mode using the F12 key.

Note: You can also edit user metadata associated with an iRule. See the example section for more information.

DISPLAY

list rule

list rule [[name] | [glob] | [regex]] ...]

show running-config rule

show running-config rule [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

show rule

show rule [[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

```
mv rule [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
options:
  to-folder
```

```
DELETE
delete rule [name]
```

```
GENERATE
generate rule [name]
options:
  checksum
  signature
```

```
HELP help rule help rule [ command | event | namespace ] [name]
```

DESCRIPTION

You cannot edit the system rules that come with the BIG-IP system. However, you can open a system rule in the editor and use it as a template to create a new rule.

To create a new rule using a system rule as a template:

1. Enter the command sequence `edit rule [system rule name]`.
tmsh opens the system rule in an editor.
2. Change the name of the rule in the editor.
3. Edit the rule and exit the editor.
tmsh checks for syntax errors, and if there are none, it saves the new rule.

For more information about iRules(r), see <http://devcentral.f5.com/>.

EXAMPLES

```
list rule
```

Displays all iRules.

```
delete rule my_irule
```

Deletes the iRule named my_irule.

```
rule my_irule {
when RULE_INIT {
}
priority 1

when SERVER_CONNECTED {
}
timing on
check strict
}
```

Creates an iRule named my_irule.

```
generate rule my_irule checksum
```

Generates a checksum for the rule definition and adds the checksum to the rule.

```
generate rule my_irule signature signing-key my_key
```

Generates a signature for the rule definition using the specified private key and adds the signature to the rule.

Note: For a rule that includes a checksum or signature to successfully load, the rule definition contents must match the stored checksum or signature. To modify the rule definition and still retain the checksum or signature, the `ignore-verification` attribute must be set to true. This is done by editing the rule and adding the `ignore-verification` attribute, which allows the modified rule to load and changes the verification status to Not Verified:

```
rule my_irule {
when RULE_INIT {}
definition-checksum 7c0dba9aa53e8959042c6cfe041d3d11
ignore-verification true }
```

Modifies an existing iRule named my_irule by adding a new metadata and modifying an existing metadata:

```
modify rule my_irule {
when RULE_INIT {}
definition-checksum 7c0dba9aa53e8959042c6cfe041d3d11
metadata replace-all-with {
  my_meta { persist false
    value "hello"
  }
  my_meta2 { persist false
    value "hello 2"
  }
} }
```

The metadata attribute is the user defined key/value pair. Metadata has the following format:

```
metadata
[add | delete | modify] {
  [metadata_name] {
    value [ "value content" ]
    persist [ true | false ]
  } }>
```

Deletes a metadata from an iRule:

```
modify rule my_irule {
when RULE_INIT {}
definition-checksum 7c0dba9aa53e8959042c6cfe041d3d11
metadata delete { my_meta } }
```

```
mv /ltm rule /Common/my_rule to-folder /Common/some_folder
```

Moves an iRule named my_rule to the folder named some_folder, where some_folder has already been created under /Common.

Note: Please note that you may not move an iRule that has an explicit usage of a configuration object, such as a pool.

OPTIONS

checksum

Generates a checksum for the rule definition and adds the checksum to the rule. This option is used only with the generate command.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the create, delete, and modify commands.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

signature

Generates a signature for the rule definition using the specified private key and adds the signature to the rule as a property. This option is used only with the generate command.

signing-key

Specifies the private key to use for signing the rule. This is used only with the signature option.

meta-data

Specifies the user-defined key/value pair associated with the rule. See the example section for usage format.

app-project

Specifies the dev plugin this rule belongs to. This is a read-only attribute.

SEE ALSO

create, delete, edit, generate, glob, list, modify, mv, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2015-10-07 ltm rule(1)

Itm snat-translation

NAME

snat-translation - Configures an explicit secure network address translation (SNAT) translation address.

MODULE

ltm

SYNTAX

Configure the snat-translation component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create snat-translation [all | [name] ]
modify snat-translation [all | [name] ]
options:
  address [ip address]
  arp [disabled | enabled]
  app-service [[string] | none]
  connection-limit [integer]
  description [string]
  [disabled | enabled]
  ip-idle-timeout [indefinite | [integer] ]
  tcp-idle-timeout [indefinite | [integer] ]
  udp-idle-timeout [indefinite | [integer] ]
  traffic-group [[string] | default | non-default | none]
```

```
edit snat-translation [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list snat-translation
list snat-translation [ [ [name] | [glob] | [regex] ] ... ]
show running-config snat-translation
show running-config snat-translation [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

```
show snat-translation
show snat-translation [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DELETE

```
delete snat-translation [all | [name] ]
```

DESCRIPTION

Explicitly defines the properties of a SNAT translation address.

EXAMPLES

```
modify snat-translation all arp disabled
```

Disables Address Resolution Protocol (ARP) on all SNAT translation addresses.

```
list snat-translation all-properties
```

Displays all properties of all SNAT translation addresses.

OPTIONS

address
The translation IP address.

arp Indicates whether the system responds to ARP requests or sends gratuitous ARPs. The default value is enabled.

app-service
Specifies the name of the application service to which this object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

connection-limit
Specifies the number of connections a translation address must reach before it no longer initiates a connection. The default value of 0 (zero) indicates that the option is disabled.

description
User defined description.

disabled
Disables SNAT translation.

enabled
Enables SNAT translation. The default value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

traffic-acceleration-status
Displays the current traffic-acceleration status. Indicates whether the SNAT address is (indirectly via a SNAT pool) in-use by or dedicated to a virtual server that uses a traffic-acceleration profile.

ip-idle-timeout
Specifies the number of seconds that IP connections initiated using a SNAT address are allowed to remain idle before being automatically disconnected. The default value is indefinite.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

tcp-idle-timeout

Specifies the number of seconds that TCP connections initiated using a SNAT address are allowed to remain idle before being automatically disconnected. The default value is indefinite.

udp-idle-timeout

Specifies the number of seconds that UDP connections initiated using a SNAT address are allowed to remain idle before being automatically disconnected. The default value is indefinite.

unit Read-only property that specifies the unit in a redundant system. Derived from traffic-group.

traffic-group

Specifies the traffic group of the failover device group on which the SNAT is active. The default traffic group is inherited from the containing folder.

inherited-traffic-group

Read-only property that indicates if the traffic-group is inherited from the parent folder.

SEE ALSO

create, delete, edit, glob, list, modify, ltm snat, ltm snatpool, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2016-10-07 ltm snat-translation(1)

ltm snat

NAME

snat - Configures secure network address translation (SNAT).

MODULE

ltm

SYNTAX

Configure the snat component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create snat [name]
modify snat [name]
options:
  (automap | none)
  auto-lasthop [default | enabled | disabled ]
  app-service [[string] | none]
  description [string]
  mirror { [disabled | enabled | none] }
  origins
    [add | delete | replace-all-with] {
[address ... | address/mask ... ]
    }
  snatpool [ name ]
  source-port [change | preserve | preserve-strict ]
  translation [translation name ... ]
  vlans
    [add | delete | replace-all-with] {
[vlan name ... ]
    }
  vlans [ default | none]
  [vlans-disabled | vlans-enabled ]
  metadata
    [add | delete | modify] {
[metadata_name ... ] {
  value [ "value content" ]
  persist [ true | false ]
}
}
```

```
edit snat [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list snat
list snat [ [ [name] | [glob] | [regex] ] ... ]
show running-config snat
show running-config snat [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

```
show snat
show snat [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  detail
  field-fmt
```

DELETE

```
delete snat [name]
```

DESCRIPTION

You can use the snat component to configure a SNAT. A SNAT defines the relationship between an externally visible IP address, SNAT IP address, or translated address, and a group of internal IP addresses, or originating addresses, of individual servers at your site.

EXAMPLES

```
create snat my_snat origins add { 10.1.1.3 } translation mySnatTranslation
```

Creates the SNAT my_snat that translates the address of connections that originate from the address 10.1.1.3 to the translation address mySnatTranslation.

```
list snat all-properties
```

Displays all properties for all SNATs.

OPTIONS

automap

Specifies that the system translates the source IP address to an available self IP address when establishing connections through the virtual server. You can use this option only if you do not use the snatpool and translation options.

Note that when you use the edit command to create a new snat, by default automap is enabled. If you do not want to use automap, you must turn this feature off by using the none option.

app-service

Specifies the name of the application service to which this object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

mirror

Enables or disables mirroring of SNAT connections. The default value is none.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

origins

Specifies a set of IP addresses and subnets from which connections originate. This option is required.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

snatpool

Specifies the name of a SNAT pool. You can only use this option if you do not use the automap and translation options.

source-port

Specifies whether the system preserves the source port of the connection. The default value is preserve.

The options are:

change

Use this setting to obfuscate internal network addresses.

preserve

Specifies to preserve the source port of the connection.

preserve-strict

Use this value only for UDP under very special circumstances such as nPath or transparent (that is, no translation of any other L3/L4 field), where there is a 1:1 relationship between virtual IP

addresses and node addresses, or when clustered multi-processing (CMP) is disabled.

translation

Specifies the name of a translated IP address. Note that translated addresses are outside the traffic management system. You can use this option only if you do not use the automap and snatpool options.

vlan

Specifies the name of the VLAN to which you want to assign the SNAT. The default value is none.

vlan-disabled

Disables the SNAT for all specified VLANs. When the "vlan" value is set to "none", the "vlan-disabled" option enables the SNAT on all VLANs.

vlan-enabled

Enables the SNAT for all specified VLANs. When the "vlan" value is set to "none", the "vlan-enabled" option disables the SNAT on all VLANs.

metadata

Associates user defined data, each of which has name and value pair and persistence. Persistent(default) means the data will be saved into config file.

SEE ALSO

create, delete, edit, glob, list, ltm snat-translation, ltm snatpool, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 ltm snat(1)

ltm snatpool

NAME

snatpool - Configures a secure network address translation (SNAT) pool.

MODULE

ltm

SYNTAX

Configure the snatpool component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

create snatpool [name]

modify snatpool [name]

options:

app-service [[string] | none]

description [string]

members

[add | delete | replace-all-with] {

[ip address ...]

}

members [default | none]

edit snatpool [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats snatpool

reset-stats snatpool [[[name] | [glob] | [regex]] ...]

DISPLAY

list snatpool

list snatpool [[[name] | [glob] | [regex]] ...]

show running-config snatpool

show running-config snatpool [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show snatpool

show snatpool [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

detail

field-fmt

DELETE
delete snatpool [name]

DESCRIPTION

A SNAT pool is a pool of translation addresses that you can map to one or more original IP addresses. Translation addresses in a SNAT pool are not self-IP addresses. You can simply create a SNAT pool and then assign it as a resource directly to a virtual server. This eliminates the need for you to explicitly define original IP addresses to which to map translation addresses.

EXAMPLES

```
create snatpool my_snat_pool1 members add { 11.12.11.24 11.12.11.25 }
```

Creates the SNAT pool my_snat_pool1 that contains the translation addresses (members) 11.12.11.24 and 11.12.11.25.

```
delete snatpool my_snat_pool1
```

Deletes the SNAT pool named my_snat_pool1.

OPTIONS

app-service

Specifies the name of the application service to which this object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

traffic-acceleration-status

Displays the current traffic-acceleration status. Indicates whether the SNAT pool is in-use by or dedicated to a virtual server that uses a traffic-acceleration profile.

members

Specifies translation IP addresses of the pools in the SNAT pool.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm snat, ltm snat-translation, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2016-10-07 ltm snatpool(1)

ltm tacdb customdb-file

NAME

customdb-file - Manages a custom TACDB file

MODULE

ltm tacdb

SYNTAX

List the customdb-file component within the ltm tacdb module using the syntax in the following sections.

CREATE

```
create customdb-file [name]
```

options:

source-path [string]

app-service [name]

DEPRECATED: Though this command is visible, this is not meant to be used by the users. The daemons use it internally.

DISPLAY

list customdb-file [[name] | all | [property]]

options:

- all-properties
- one-line
- partition

DELETE

delete customdb-file [[name] | all]

DEPRECATED: Though this command is visible, this is not meant to be used by the users. The daemons use it internally.

DESCRIPTION

The customdb-file is created internally using the URI specified in `ltm tacdb customdb`.

EXAMPLES

```
create customdb-file new { source-path file:/shared/images/new.txt }
```

Creates a new customdb file object, "new" from the source file `/shared/images/new.txt`

```
list customdb-file [fileobj-name]
```

Lists the attributes of customdb file object, "new" from the source file `/shared/images/new.txt`

OPTIONS

create

Creates a new file object for custom tacdb.

delete

Deletes the file object that you specify next.

source-path

Specifies the location from where the TACDB file object sources the file.

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the `strict-updates` option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description

User defined description for this customdb file object.

partition

Displays the administrative partition within which the component resides.

SEE ALSO

`edit`, `list`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2015. All rights reserved.

BIG-IP 2015-05-12 ltm tacdb customdb-file(1)

ltm tacdb customdb

NAME

customdb - Configures a custom tacdb to be used for file loads. A customdb includes a URL (including local file paths) from where TACDB files are loaded. These files contain TAC code information.

MODULE

ltm tacdb

SYNTAX

Configure the customdb component within the ltm tacdb module using the syntax in the following sections.

CREATE/MODIFY

```
create customdb [name]
```

```
modify customdb [[name] | all]
```

options:

url [string]

poll-interval [integer]

user [string]

password [string]

priority [high | low]
app-service [name]
description [string]

DISPLAY

list customdb [[name] | all | [property]]

options:

all-properties
non-default-properties
one-line
partition
recursive

DELETE

delete customdb [[name] | all]

LOAD load customdb [name]

DESCRIPTION

You can use the customdb component to define a custom tacdb

EXAMPLES

```
create customdb new { url file:/shared/images/new.txt }
```

Creates a new custom tacdb, "new" with TAC code information in the file specified by url.

```
delete customdb new
```

Deletes the custom tacdb that has been specified.

```
load customdb new
```

Trigger a manual download of the custom tacdb from the location pointed to and load it into memory.

OPTIONS

url Specifies the url to fetch the file containing the TACDB information.

poll-interval

Specifies the time interval in seconds at which the URL needs to be polled. The initial default value 0 disables polling. The literal "default" value is 300 seconds (5 minutes).

priority

Specifies the priority whether the customdb is higher or lower than built-in/licensed TACDB.

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description

User defined description for this customdb.

partition

Displays the administrative partition within which the component resides.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

BIG-IP 2017-05-17 Itm tacdb customdb(1)

Itm tacdb licenseddb

NAME

licenseddb - Configures a licensed tacdb to be used for file loads.

MODULE

itm tacdb

SYNTAX

Configure the licenseddb component within the Itm tacdb module using the syntax in the following sections.

MODIFY

modify licenseddb [licensed-tacdb | all]

options:
poll-interval [integer]

LOAD
load licenseddb [licensed-tacdb]

DISPLAY
list licenseddb [licensed-tacdb | all | [property]]
options:
all-properties
non-default-properties
one-line
recursive

DESCRIPTION
You can use the licenseddb component to define polling interval of downloading licensed tacdb

EXAMPLES
modify licenseddb licensed-tacdb poll-interval 2

Modifies polling interval, "2" means two weeks.

load licenseddb licensed-tacdb

Triggers a manual download of the licensed tac-db from the built-in location in the F5 servers and loads it into memory.

OPTIONS
poll-interval
Specifies the time interval in weeks at which the licensed tacdb needs to be polled

SEE ALSO
edit, list, modify, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2015. All rights reserved.

BIG-IP 2017-05-17 ltm tacdb licenseddb(1)

ltm tacdb query

NAME
query - Displays the device type and OS to which a TAC code belongs

MODULE
ltm tacdb

SYNTAX
Use the query component within the ltm tacdb module to query the device type and OS the TAC code belongs to.

DISPLAY
show ltm tacdb query [integer]

DESCRIPTION
You can use the query component to query the device type and OS a url belongs to

EXAMPLES
show ltm tacdb query 35333303
Displays the device type and OS "35333303" belongs to.

OPTIONS
integer
Specify the TAC code.

SEE ALSO
show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

ltm traffic-class

NAME

traffic-class - Configures a traffic class.

MODULE

ltm

SYNTAX

Configure the traffic-class component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create traffic-class [name]
```

```
modify traffic-class [name]
```

options:

```
app-service [[string] | none]
```

```
classification [string]
```

```
description [string]
```

```
destination-address [ [ip address] | none]
```

```
destination-mask [ [ip address] | none]
```

```
destination-port [ [integer] | [port name] ]
```

```
protocol [any | [protocol] ]
```

```
source-address [ [ip address] | none]
```

```
source-mask [ [ip address] | none]
```

```
source-port [ [integer] | [port name] ]
```

```
edit traffic-class [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DISPLAY

```
list traffic-class
```

```
list traffic-class [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config traffic-class
```

```
show running-config traffic-class [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DELETE

```
delete traffic-class [name]
```

DESCRIPTION

You can use the traffic-class component to configure a traffic class, which is a named group of ports, machines, and subnets. You can then assign this traffic class to a virtual server in order to configure the virtual server to achieve specific Quality of Service (QoS) standards.

EXAMPLES

```
create traffic-class my_traffic_class classification "My traffic class."
```

Creates a traffic class named my_traffic_class, which tags matching flows with the tag My traffic class.

```
list traffic-class all-properties
```

Displays all of the properties of all of the traffic classes.

```
delete traffic-class my_traffic_class
```

Deletes the traffic class named my_traffic_class.

OPTIONS

app-service

Specifies the name of the application service to which the traffic class belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the traffic class. Only the application service can modify or delete the traffic class.

classification

Specifies the actual textual tag to be associated with the flow if the traffic class is matched. This option is required.

description

User defined description.

destination-address

Specifies destination IP addresses for the system to use when evaluating traffic flow. If traffic flow matches this value, it is tagged with the value in the classification option. The default value is none.

destination-mask

Specifies a destination IP address mask for the system to use when evaluating traffic flow. If traffic flow matches this value, it is tagged with the value in the classification option. The default value is none.

destination-port

Specifies a destination port for the system to use when evaluating traffic flow. If traffic flow matches this value, it is tagged with the value in the classification option. The default value is 0 (zero).

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

protocol

Specifies a protocol for the system to use when evaluating traffic flow. If traffic flow matches this value, it is tagged with the value in the classification option. The default value is any.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

source-address

Specifies source IP addresses for the system to use when evaluating traffic flow. If traffic flow matches this value, it is tagged with the value in the classification option. The default value is none.

source-mask

Specifies a source IP address mask for the system to use when evaluating traffic flow. If traffic flow matches this value, it is tagged with the value in the classification option. The default value is none.

source-port

Specifies a source port for the system to use when evaluating traffic flow. If traffic flow matches this value, it is tagged with the value in the classification option. The default value is 0 (zero).

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 ltm traffic-class(1)

ltm traffic-matching-criteria

NAME

traffic-matching-criteria - Configures a virtual server matching object that defines how traffic is steered towards the virtual server object for the Local Traffic Manager.

MODULE

ltm

SYNTAX

Configure the traffic-matching-criteria component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

create traffic-matching-criteria [name]
modify traffic-matching-criteria [name]

options:

all

description [string]

destination-address-list [[addr_list_name] | none]

destination-address-inline [[address] | none]

destination-port-list [[port_list_name] | none]

destination-port-inline [port]

source-address-list [[addr_list_name] | none]

source-address-inline [address]

source-port-inline [port]

edit traffic-matching-criteria [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties

DISPLAY

list traffic-matching-criteria

list traffic-matching-criteria [[[name] | [glob] | [regex]] ...]

show running-config traffic-matching-criteria

show running-config traffic-matching-criteria [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line
- partition

show traffic-matching-criteria

show traffic-matching-criteria [name]

options:

- hours
- minutes
- seconds

DELETE

delete traffic-matching-criteria [name]

Note: You must remove all references to a traffic-matching-criteria before you can delete the traffic-matching-criteria.

DESCRIPTION

You can use this traffic-matching-criteria component to configure the virtual server matching definitions on the Local Traffic Manager. This defines the criteria used for matching destination and source address/ports for a virtual server.

EXAMPLES

```
create traffic-matching-criteria my_traffic_matching_criteria destination-address-list dst_addr_list
destination-port-list dst_port_list
```

Creates a Local Traffic Manager traffic-matching-criteria named my_traffic_matching_criteria with dst_addr_list and the dst_port_list for destination matching, from any source.

```
create traffic-matching-criteria my_traffic_matching_criteria source-address-list src_addr_list source-port-
list src_port_list
```

Creates a Local Traffic Manager traffic-matching-criteria named my_traffic_matching_criteria with src_addr_list and the port-list for source matching, to any destination.

```
create traffic-matching-criteria my_traffic_matching_criteria destination-address-list dst_addr_list
destination-port-list dst_port_list source-address-list src_addr_list1 source-port-list src_port_list1
```

Creates a Local Traffic Manager traffic-matching-criteria named my_traffic_matching_criteria with src_addr_list and the src_port-list for source matching, and dst_addr_list and dst_port_list for destination matching.

```
delete traffic-matching-criteria my_traffic_matching_criteria
```

Deletes the traffic_matching_criteria named my_traffic_matching_criteria.

```
show traffic-matching-criteria
```

Displays status for all Local Traffic Manager traffic-matching-criteria objects in the system configuration.

```
show traffic-matching-criteria all-properties
```

Displays status for all Local Traffic Manager traffic-matching-criteria objects in the system configuration.

```
list traffic-matching-criteria my_traffic_matching_criteria
```

Displays properties of the traffic-matching-criteria named my_traffic_matching_criteria.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

destination-address-list

Specifies the name of destination address list for this virtual server matching object i.e. the addresses on which the virtual server referencing this object will listen for connections. The default value is any.

destination-address-inline

Specifies a single destination address for this virtual server matching object i.e. the address on which the virtual server referencing this object will listen for connections. The default value is none.

destination-port-list

Specifies the name of destination port list for this virtual server matching object i.e. the ports on which the virtual server referencing this object will listen for connections. The default value is any.

destination-port-inline

Specifies a single destination port for this virtual server matching object i.e. the port on which the virtual server referencing this object will listen for connections. The default value is none.

source-address-list

Specifies the source address list for this virtual server matching object i.e. the addresses from which the virtual server referencing this object will accept traffic. The default value is any.

source-address-inline

Specifies a single source address list for this virtual server matching object i.e. the address from which the virtual server referencing this object will accept traffic. The default value is none.

source-port-list

Specifies the source port list for this virtual server matching object i.e. the ports from which the virtual server referencing this object will accept traffic. The default value is any.

source-port-inline

Specifies a single source port for this virtual server matching object i.e. the port from which the virtual server referencing this object will accept traffic. The default value is none.

description

User defined description.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, net ip-address-list, net port-list, ltm virtual, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2015, 2017. All rights reserved.

BIG-IP 2017-09-05 ltm traffic-matching-criteria(1)

ltm urlcat-cloud-cache

NAME

urlcat-cloud-cache - Operating on the cache of cloud lookup result.

MODULE

ltm

SYNTAX

Use the urlcat-cloud-cache component within the ltm module to operating on the cloud cache.

DISPLAY

delete ltm urlcat-cloud-cache url all

DESCRIPTION

You can use the urlcat-cloud-cache component to operating on the cache.

EXAMPLES

delete ltm urlcat-cloud-cache url all

Delete all the cache entries currently in cloud cache.

OPTIONS

all Use all to delete all the cache entries currently in cloud cache.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2015-09-01 ltm urlcat-cloud-cache(1)

Itm urlcat-query

NAME

urlcat-query - Displays the category to which a url belongs

MODULE

itm

SYNTAX

Use the urlcat-query component within the Itm module to query the category the url belongs to.

DISPLAY

```
show Itm urlcat-query [string]
```

DESCRIPTION

You can use the urlcat-query component to query the category a url belongs to

EXAMPLES

```
show Itm urlcat-query www.google.com
```

Displays the category "www.google.com" belongs to.

OPTIONS

string

Specify the urlname.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014. All rights reserved.

BIG-IP 2014-12-26 Itm urlcat-query(1)

Itm virtual-address

NAME

virtual-address - Configures virtual addresses.

MODULE

itm

SYNTAX

Configure the virtual-address component within the Itm module using the syntax shown in the following sections.

CREATE/MODIFY

```
create virtual address [name]
```

```
modify virtual address [name]
```

options:

```
address [ip address]
```

```
app-service [[string] | none]
```

```
arp [enabled | disabled]
```

```
auto-delete [true | false]
```

```
connection-limit [integer]
```

```
description [string]
```

```
enabled [yes | no]
```

```
icmp-echo [enabled | disabled | selective | always | any | all]
```

```
mask [netmask]
```

```
route-advertisement [enabled | disabled | selective | always | any | all]
```

```
server-scope [all | any | none]
```

```
spanning [enabled | disabled]
```

```
traffic-group [[string] | default | non-default | none]
```

metadata

```
[add | delete | modify] {
```

```
[metadata_name ... ] {
```

```
value [ "value content" ]
```

```
persist [ true | false ]
```

```
}
```

```
}
```

```
edit virtual-address [ [ [name] | [glob] | [regex] ] ... ]
```

options:
all-properties
non-default-properties

reset-stats virtual-address
reset-stats virtual-address [[[name] | [glob] | [regex]] ...]

mv virtual-address [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]
options:
to-folder

DISPLAY
list virtual-address
list virtual-address [[[name] | [glob] | [regex]] ...]
show running-config virtual-address
show running-config virtual-address
[[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

show virtual-address
show virtual-address [[[name] | [glob] | [regex]] ...]
options:
all-properties
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

DELETE
delete virtual-address [all | [name]]

DESCRIPTION

You can use the virtual-address component to enable, disable, display, and delete virtual addresses. You can also list the virtual address configuration, and view statistics for a specific virtual address.

Note that tmsh only displays virtual addresses when you explicitly request them. For example:

To display the properties of virtual addresses or a specific virtual address from the ltm module, use the command sequences list virtual-address and list virtual-address [name], respectively.

To display statistics for virtual addresses or a specific virtual address from the ltm module, use the command sequence show virtual-address and show virtual-address [name], respectively.

EXAMPLES

```
create virtual-address myVirtualAddr address 10.10.10.20 enabled yes
```

Creates a virtual address 10.10.10.20, with a name of myVirtualAddr.

```
create virtual-address myVirtualAddr address 10.10.10.20 enabled yes traffic-group /Common/traffic-group-1
```

Creates a virtual address 10.10.10.20, with a name of myVirtualAddr, that is assigned to traffic-group-1.

```
modify virtual-address myVirtualAddr enabled no
```

Disables the virtual address myVirtualAddr.

```
delete virtual-address myVirtualAddr
```

Deletes the virtual address myVirtualAddr.

```
list virtual-address myVirtualAddr all-properties
```

Lists the configuration information for the virtual address, myVirtualAddr.

```
show virtual-address myVirtualAddr
```

Displays statistics and status for the virtual-address myVirtualAddr.

```
show virtual-address myVirtualAddr all-properties
```

Displays statistics and status for the virtual named myVirtualAddr.

Note that if the system includes Packet Velocity(r) ASIC (PVA) and PVA Assist capabilities, this command displays status and statistics for that feature.

```
mv /ltm virtual-address /Common/10.10.10.20 to-folder /Common/all_virtual_addresses
```

Moves the virtual-address 10.10.10.20 to a folder named all_virtual_addresses.

Note: If you wish to change the name of the virtual-address, you must use the configured IP Address or a name that does not represent a different IP Address than the one configured.

Please refer to the mv manual page for additional examples on how to use the mv command.

OPTIONS

address

The virtual IP address.

arp Enables or disables ARP for the specified virtual address. The default value is enabled.

app-service

Specifies the name of the application service to which the virtual address belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the virtual address. Only the application service can modify or delete the virtual address.

auto-delete

Indicates if the virtual address will be deleted automatically on deletion of the last associated virtual server or not. The default value is true.

connection-limit

Sets a concurrent connection limit for one or more virtual servers. The default value is 0, meaning "no limit."

description

User defined description.

enabled

Specifies whether the specified virtual address is enabled. The default value is yes.

floating

Read-only property derived from traffic-group. A floating virtual address is a virtual address for a VLAN that serves as a shared address by all devices of a BIG-IP traffic-group.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

icmp-echo

Specifies whether the virtual address should reply to ICMP echo requests. The default value is enabled.

enabled or always

ICMP echo reply will be sent in response to ICMP echo requests.

disabled

ICMP echo reply will not be sent in response to ICMP echo requests.

selective

ICMP echo reply will be sent in response to ICMP echo requests, when availability status is true.

any ICMP echo reply will be sent in response to ICMP echo requests, when any of the contributing virtual server is available.

all ICMP echo reply will be sent in response to ICMP echo requests when all of the contributing virtual server is available.

mask Sets the netmask for one or more network virtual servers only. This setting is required for network virtual servers. The default value is 255.255.255.255.

partition

Displays the administrative partition within which the virtual address resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

route-advertisement

Specifies the route advertisement setting for the virtual address. The default value is disabled.

enabled or selective

Route is advertised when virtual-address is available.

disabled

Route is not advertised.

always

Route is advertised regardless of the availability status.

any Route is advertised when any of the contributing virtual server is available.

all Route is advertised when all of the contributing virtual server is available.

server-scope

Specifies when the virtual address is considered available. When a virtual address is available and Route Advertisement is enabled or selective, the BIG-IP system advertises the route for the virtual address. The default value is any.

any When any virtual server is available: Advertises the route when any virtual server is available.

all When all virtual servers are available: Advertises the route when all virtual servers are available.

none Always advertises the route regardless of the virtual servers available.

spanning

Enables or disables spanning for the specified virtual address. The default value is disabled.

unit Read-only property that specifies the unit in a redundant system. Based on traffic-group.

traffic-group

Specifies the traffic group on which the virtual address is active. The default traffic group is inherited from the containing folder.

inherited-traffic-group

Read-only property that indicates if the traffic-group is inherited from the parent folder.

metadata

Associates user defined data, each of which has name and value pair and persistence. Persistent(default) means the data will be saved into config file.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, mv, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2018-03-06 ltm virtual-address(1)

ltm virtual

NAME

virtual - Configures a virtual server.

MODULE

ltm

SYNTAX

Configure the virtual component within the ltm module using the syntax shown in the following sections.

CREATE/MODIFY

create virtual [name]

modify virtual [name]

options:

all

address-status [yes | no]

app-service [[string] | none]

auth [add | delete | replace-all-with] {
[profile_name ...]

}

auth [default | none]

auto-discovery [enabled | disabled]

auto-lasthop [default | enabled | disabled]

clone-pools [add | delete | replace-all-with] {
[pool_name ...] {

}

context [clientside | serverside]

}

clone-pools none

cmp-enabled [yes | no]

connection-limit [integer]

dhcp-relay

description [string]

destination [[virtual_address_name:port] | [ipv4:port] | [ipv6.port]]

[disabled | enabled]

eviction-protected [enabled | disabled]

fallback-persistence [none | [profile name]]

flow-eviction-policy [none | [eviction policy name]]

fw-enforced-policy [[policy_name] | none]

fw-staged-policy [[policy_name] | none]

gtm-score [integer]

ip-forward

ip-protocol [any | [protocol]]

internal

l2-forward

last-hop-pool [[pool_name] | none]

mask { [ipv4] | [ipv6] }

mirror { [disabled | enabled | none] }

```

nat64 [enabled | disabled]
per-flow-request-access-policy [ [policy_name] | none ]
persist [replace-all-with] {
  [profile_name ... ] {
default [no | yes]
  }
}
persist none
policies [ add | delete | replace-all-with] {
  policy_name [[policy_name] ...]
}
pool [ [pool_name] | none]
profiles [add | delete | replace-all-with] {
  [profile_name ...] {
context [all | clientside | serverside]
  }
}
profiles [default | none]
rate-class [name]
rate-limit [integer]
rate-limit-mode [destination | object | object-destination |
  object-source | object-source-destination | source |
  source-destination]
rate-limit-dst [integer]
rate-limit-src [integer]
related-rules { none | [rule_name ...] }
reject
rules { [none | [rule_name ...] ] }
security-nat-policy {
  policy [ [policy_name] | none]
  use-device-policy [no | yes]
  use-route-domain-policy [no | yes]
}
serverssl-use-sni [ enabled | disabled ]
service-down-immediate-action [none | drop | reset]
service-policy [ [policy_name] | none ]
snat [automap | none] DEPRECATED - see source-address-translation
snatpool [snatpool_name] DEPRECATED - see source-address-translation
source { [ipv4[/prefixlen]] | [ipv6[/prefixlen]] }
source-address-translation {
  options:
    pool [ [pool_name] | none]
    type [ automap | lsn | snat | none ]
}
source-port [change | preserve | preserve-strict]
traffic-classes [add | delete | replace-all-with] {
  [traffic_class_name ...]
}
traffic-classes [default | none]
translate-address [enabled | disabled]
translate-port [enabled | disabled]
transparent-nextthop [vlan_name]
vlans [add | delete | replace-all-with] {
  [vlan_name ... ]
}
vlans [default | none]
vlans-disabled
vlans-enabled
metadata [add | delete | modify] {
  [metadata_name ... ] {
value [ "value content" ]
persist [ true | false ]
  }
}
reset-stats virtual [ [ [name] | [glob] | [regex] ] ... ]
fw-enforced-policy-rules { [rule name] }
fw-staged-policy-rules { [rule name] }
security-nat-rules { [rule name] }
profiles { [profile name] }

options:
  fw-context-stat
  ip-intelligence-categories
  port-misuse

DISPLAY
list virtual
list virtual [ [ [name] | [glob] | [regex] ] ...]
show running-config virtual
show running-config virtual [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

show virtual

```

show virtual [[[name] | [glob] | [regex]] ...]

options:

all-properties (default | exa | gig | kil | meg | peta | raw | tera |
yotta | zetta)
detail
field-fmt
fw-context-stat
ip-intelligence-categories
port-misuse

mv virtual [[[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]]]

options:

to-folder

DELETE

delete virtual [name]

DESCRIPTION

You can use the virtual component to create, delete, modify properties on, and display information about virtual servers. Virtual servers are externally visible IP addresses that receive client requests. Rather than sending the requests directly to the destination IP address specified in the packet header, it sends the requests to any of several content servers that make up a load balancing pool. Virtual servers also apply various behavioral settings to multiple traffic types, enable persistence for multiple traffic types, and direct traffic according to user-written iRules(r).

Note: After you configure a Global Traffic Manager listener, when you use the tab completion feature within the ltm module, the listener displays as one of the virtual servers in the Configuration Items section.

EXAMPLES

```
create virtual myV2 { destination 11.11.11.12:any persist replace-all-with { source_addr } } pool myPool}
```

Creates a virtual server named myV2, which uses the source address persistence method.

```
modify virtual vs_f14_http4 profiles replace-all-with { profile-udp }
```

Replaces the profile associated with the virtual server vs_f14_http4.

Note: To replace the profile associated with a virtual server, you must enclose the name of the new profile in curly brackets.

```
delete virtual myV4 myV5 myV6
```

Deletes the virtual servers named myV4, myV5, and myV6.

```
show virtual myV4
```

Displays statistics and status for the virtual named myV4.

```
show virtual myV4 all-properties
```

Displays statistics and status for the virtual named myV4.

Note: If the system includes Packet Velocity(r) ASIC (PVA) and PVA Assist capabilities, this command displays status and statistics for that feature.

```
mv /ltm virtual /Common/my_vip to-folder /Common/some_folder
```

Moves a virtual server named my_vip to the folder named some_folder, where some_folder has already been created under /Common.

Note: Please note that you may not move a virtual server that is associated with CGNAT configuration items, such as LSN pools.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

address-status

Specifies whether the virtual will contribute to the operational status of the associated virtual-address. The default value is 'yes'.

app-service

Specifies the name of the application service to which the virtual server belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the virtual server. Only the application service can modify or delete the virtual server.

auth Specifies a list of authentication profile names, separated by spaces, that the virtual server uses to manage authentication.

auto-discovery

Enable or disable security protected objects (virtual server) auto discovery functionality. The default value is disabled.

clone-pools

Specifies a pool or list of pools that the virtual server uses to replicate either client or server traffic. You must specify a value of either clientside or serverside for the context option for each clone pool. Typically, this option is used for intrusion detection.

cmp-enabled

Enables or disables clustered multi-processor (CMP) acceleration. This feature applies to certain platforms only. The default value is yes.

connection-limit

Specifies the maximum number of concurrent connections you want to allow for the virtual server. The default value of 0 (zero) allows for an unlimited number of concurrent connections.

context

Specifies that the pool is either a clientside or serverside clone pool.

Note: Because validation occurs outside of TMSH, you will receive an error when you modify the context for profiles in a virtual server.

dhcp-relay

Specifies a virtual server that relays all received dhcp requests to all pool members. If there is no pool, the received request get dropped. If you specify the dhcp-relay option, you cannot use the ip-forward or l2-forward or reject options.

description

User defined description.

destination

Specifies the name of the virtual address and service on which the virtual server listens for connections.

The format for "ipv4" is a.b.c.d[:port]. The format for an "ipv6" address is a:b:c:d:e:f:g:h[:port].

The default value is any:any.

(enabled | disabled)

Specifies the state of the virtual server. The default value is enabled.

Note: When you disable a virtual server, the virtual server no longer accepts new connection requests. However, it allows current connections to finish processing before going to a down state.

eviction-protected

Enables or disables protection for the virtual server from the aggressive sweeper. The default is disabled.

fallback-persistence

Specifies a fallback persistence profile for the virtual server to use when the default persistence profile is not available. The default value is none.

flow-eviction-policy

Specifies a flow eviction policy for the virtual server to use, to select which flows to terminate when the number of connections approaches the connection limit on the virtual server. The default value is none.

fw-enforced-policy

Specifies an enforced firewall policy. fw-enforced-policy rules are enforced on a virtual server.

fw-enforced-policy-rules

Specifies firewall rules enforced on ltm virtual via referenced fw-enforced-policy.

fw-staged-policy

Specifies a staged firewall policy. fw-staged-policy rules are not enforced while all the visibility aspects namely statistics, reporting and logging function as if the fw-staged-policy rules were enforced on a virtual server.

fw-staged-policy-rules

Specifies firewall rules staged on ltm virtual via referenced fw-staged-policy.

security-nat-rules

Specifies security nat rules associated with ltm virtual via referenced security-nat-policy.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

gtm-score

Specifies a score that is associated with the virtual server. Global Traffic Manager (GTM) can rely on this value to load balance traffic in a proportional manner.

traffic-acceleration-status

Displays the current traffic-acceleration status. The virtual server is considered traffic-acceleration-dedicated if it uses a traffic-acceleration profile.

ip-forward

Specifies a virtual server that has no pool members to load balance, but instead, forwards the packet directly to the destination IP address specified in the client request. If you specify the ip-forward option, you cannot use the l2-forward or reject options. The destination, mask, translate-address, translate-port, vlans, vlans-disabled and vlans-enabled attributes are set by the system, any attempt to change them will have no effect.

ip-protocol

Specifies the IP protocol for which you want the virtual server to direct traffic. Sample protocol names

are TCP and UDP. The default value is any.

Note: You do not use this setting when creating an HTTP class virtual server.

internal

Specifies an internal virtual server that handles requests for a parent virtual server, such as content adaptation. Internal virtual servers do not receive external connections, instead they are specified by name by profiles in the parent virtual server (see ltm profile request-adapt and ltm profile response-adapt). Since internal virtual servers do not listen for external connections, not all attributes are used for internal virtual servers. The destination, mask, translate-address, translate-port, vlans, vlans-disabled and vlans-enabled attributes are set by the system, any attempt to change them will have no effect.

l2-forward

Specifies a virtual server that shares the same IP address as a node in an associated VLAN. You create this type of virtual server when you want to create a VLAN group. If you specify the l2-forward option, you cannot use the ip-forward or reject options.

last-hop-pool

Specifies the name of the last hop pool that you want the virtual server to use to direct reply traffic to the last hop router. The default value is none.

mask Specifies the netmask for a network virtual server only. This setting is required for a network virtual server.

The netmask clarifies whether the host bit is an actual zero or a wildcard representation. The default value is 255.255.255.255 for IPv4 or ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff for IPv6.

mirror

Enables or disables mirroring. You can use mirroring to maintain the same state information in the standby unit that is in the active unit, allowing transactions such as FTP file transfers to continue as though uninterrupted. The default value is none.

mobile-app-tunnel

Deprecated since v13.1.0.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

nat64

Enable or disable NAT64. The default value is disabled. NAT64 is a service that automatically translate IPv6 traffic into IPv4.

partition

Displays the name of the administrative partition within which the virtual server resides.

per-flow-request-access-policy

Specifies the name of the per-request access policy to be used with the virtual server. The default value is none.

persist

Specifies a list of profiles separated by spaces that the virtual server uses to manage connection persistence. The default value is none.

To enable persistence, typically you specify a single profile. However, you can specify multiple profiles in conjunction with iRules(r) that define a persistence strategy based on incoming traffic. In the case of multiple profiles, the default option specifies which profile you want the virtual server to use if an iRule does not specify a persistence method. When you specify multiple profiles, the default value of the default property is no. You can set the value of the default property to yes for only one of the profiles.

policies

Manage LTM Policies applied to the virtual server. LTM Policies define a set of conditions and actions that can be used to inspect, modify, direct traffic, and enable/disable features on the fly, similar to iRules. LTM Policies do not require programming. See also ltm policy.

pool Specifies a default pool to which you want the virtual server to automatically direct traffic. The default value is none.

port-misuse

Used to show or reset port misuse policy statistics for the virtual server.

fw-context-stat

Used to show or reset firewall statistics for the virtual server.

profiles

Specifies a list of profiles for the virtual server to use to direct and manage traffic. The default value is fastL4.

rate-class

Specifies the name of an existing rate class that you want the virtual server to use to enforce a throughput policy for incoming network traffic. The default value is none.

rate-limit

Specifies the maximum number of connections per second allowed for a virtual server. The default value is 'disabled'.

rate-limit-mode

Indicates whether the rate limit is applied per virtual object, per source address, per destination address, or some combination thereof. The default value is 'object', which does not use the source or destination address as part of the key.

rate-limit-dst-mask

Specifies a mask, in bits, to be applied to the destination address as part of the rate limiting. The default value is '0', which is equivalent to using the entire address - '32' in IPv4, or '128' in IPv6.

rate-limit-src-mask

Specifies a mask, in bits, to be applied to the source address as part of the rate limiting. The default value is '0', which is equivalent to using the entire address - '32' in IPv4, or '128' in IPv6.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

related-rules

Specifies a list of iRules, separated by spaces, that customize the behavior of secondary channels (for instance the data channel on FTP) opened on behalf of the virtual server. The default value is none.

reject

Specifies that the BIG-IP(r) system rejects any traffic destined for the virtual server IP address. If you specify the reject option, you cannot use the ip-forward or l2-forward options.

rules

Specifies a list of iRules, separated by spaces, that customize the virtual server to direct and manage traffic. The default value is none.

security-nat-policy

Configures the following options to specify which Security NAT Policy is to be used to match the incoming traffic and perform source/destination translation (address/port) using the first-match rule criteria:

policy

Specifies the name of the Security NAT Policy to be used (see security nat policy).

use-route-domain-policy

Specifies whether to use the virtual server's route domain context's Security NAT policy. If enabled AND the virtual server does not have a NAT policy configured, route domain's security NAT policy is used.

use-device-policy

Specifies whether to use the security device context NAT policy (see security device-context). If enabled AND both virtual server as well as route domain do not have a NAT policy configured, NAT policy configured at security device (a.k.a global) level is used.

serverssl-use-sni

When multiple server-ssl profiles are attached to a virtual, setting this allows one to be chosen based on the SNI extension from the ClientHello if a client-ssl profile is also attached to the virtual.

service-down-immediate-action

Specifies the immediate action the BIG-IP system should respond with upon the receipt of the initial client's SYN packet if the availability status of the virtual server is Offline or Unavailable. This is supported for the virtual server of Standard type and TCP protocol. The default value is none.

service-policy

Specifies a service policy for the virtual server. If set, it will enforce the service policy for incoming network traffic. The service policy can be used to validate if incoming traffic conforms to a set of application protocols.

snat Specifies whether SNAT automap is enabled for the virtual server. The default value is none. This attribute is DEPRECATED. Use source-address-translation { type (automap / none) }

snatpool

Specifies the name of an existing SNAT pool that you want the virtual server to use to implement selective and intelligent SNATs. This attribute is DEPRECATED. Use source-address-translation { type snatpool pool pool_name }

source

Specifies an IP address or network from which the virtual server will accept traffic.

The format for an "ipv4" address is a.b.c.d[/prefixlen]. The format for an "ipv6" address is a:b:c:d:e:f:g:h[/prefixlen].

source-address-translation

Specifies the type of source address translation enabled for the virtual server as well as the pool that the source address translation will use.

pool Specifies the name of a LSN or SNAT pool used by the specified virtual server.

type Specifies the type of source address translation associated with the specified virtual server.

The options are:

automap

Specifies the use of self IP addresses for virtual server source address translation.

`lsn` Specifies the use of a LSN pool of translation addresses for virtual server source address translation.

`none` Specifies no source address translation to be used by the virtual server.

`snat` Specifies the use of a SNAT pool of translation addresses for virtual server source address translation.

`source-port`

Specifies whether the system preserves the source port of the connection. The default value is `preserve`.

The options are:

`change`

Obfuscates internal network addresses.

`preserve`

Preserves the source port of the connection.

`preserve-strict`

Use this value only for UDP under very special circumstances, such as `nPath` or `transparent` (that is, no translation of any other L3/L4 field), where there is a 1:1 relationship between virtual IP addresses and node addresses, or when clustered multi-processing (CMP) is disabled.

`traffic-classes`

Specifies a list of traffic classes that are associated with the virtual server. The default value is `none`.

`translate-address`

Enables or disables address translation for the virtual server. Disable address translation for a virtual server if you want to use the virtual server to load balance connections to any address. This option is useful when the system is load balancing devices that have the same IP address. The default value is disabled.

`translate-port`

Enables or disables port translation. Disable port translation for a virtual server, if you want to use the virtual server to load balance connections to any service. The default value is disabled.

`transparent-nexthop`

Specifies the egress interface for traffic and enables layer 2 (MAC) address preservation. Layer 2 address preservation disables layer 3 (IP/IPv6) address translation.

`vlan`

Specifies a list of VLANs on which the virtual server is either enabled or disabled. The default value is `none`. The options `vlan-disabled` and `vlan-enabled` indicate whether the virtual server is disabled or enabled on the list of specified VLANs.

`vlan-disabled`

Disables the virtual server on the VLANs specified in the `vlan` option. This is the default setting.

`vlan-enabled`

Enables the virtual server on the VLANs specified in the `vlan` option.

`vs-index`

Displays a unique index assigned to this virtual server.

`metadata`

Associates user defined data, each of which has name and value pair and persistence. `Persistent`(default) means the data will be saved into config file.

`ip-intelligence-categories`

Used to show/ reset statistics on IP intelligence white/ black lists categories.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltn persistence`, `ltn pool`, `modify`, `mv`, `security nat policy`, `net service-policy`, `net vlan`, `net vlan-group`, `security firewall schedule`, `security firewall rule-list`, `regex`, `reset-stats`, `rule`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2019-06-19 ltm virtual(1)

mgmt shared settings api-status availability

NAME

availability - Configures the accessibility of the elements that have one of the specified API lifecycle states.

MODULE

mgmt shared settings api-status

SYNTAX

Configure the availability component within the mgmt shared settings api-status module using the syntax shown in the following sections.

MODIFY

modify availability

options:

deprecatedApiAllowed [true | false]

earlyAccessApiAllowed [true | false]

testOnlyApiAllowed [true | false]

DISPLAY

list availability

options:

deprecatedApiAllowed

earlyAccessApiAllowed

testOnlyApiAllowed

DESCRIPTION

There are six F5 API lifecycle states namely: internal, testOnly, earlyAccess, generalAccess, deprecated, noStatus. These states are described at the end of this page. For the commands, command properties in deprecated, earlyAccess and testOnly states the availability can be configured by the user.

The availability settings indicate if the commands or properties within a command are accessible to the user. For example: if the deprecatedApiAllowed property is set to false; then it implies that the commands and properties with deprecated status are not available to user. i.e. These elements will not be available for create, modify, delete commands. For list action, the elements will not displayed at the command line.

Warnings will be generated at stderr and /var/log/ltm based on the mgmt::shared::settings::api-status::log::resource and resource-property settings.

If these settings are modified in the tmsh interactive mode then, the user has to exit the tmsh interactive mode and re-enter for the settings to be effective.

OPTIONS

deprecatedApiAllowed

Determines the accessibility of commands, properties in deprecated state. The default value is true.

earlyAccessApiAllowed

Determines the accessibility of commands, properties in earlyAccess state. The default value is true.

testOnlyApiAllowed

Determines the accessibility of commands, properties in testOnly state. The default value is false.

The remaining properties are read-only and are not user modifiable.

F5 API lifecycle policy

The F5 API lifecycle consists of five states, namely: internal, testOnly, earlyAccess, generalAccess, deprecated, noStatus. These states can be described as follows:

internal

Internal routing not exposed through public interface

testOnly

Enabled only when test variable configured. Used internally for test cases that require special-case workers to complete functional testing for features not exposed in public API.

earlyAccess

Experimental and susceptible to change in future releases. Used for new features that haven't had time to be solidified.

generalAccess

API that satisfies the API general access release policy.

deprecated

Resource still exists in the API but indicates that there is an preferred alternative or that it may be removed in the future based on deprecation policy.

noStatus

For backward compatibility for APIs not categorized. At some point this status will be disallowed and build time break will occur if resource not categorized with appropriate flag.

In the write-up above the terms command, command::properties are used in place of API to describe the behavior of the mgmt::shared::settings::api-status::availability settings.

SEE ALSO

mgmt shared settings api-status log resource, mgmt shared setting api-status log resource-property

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2017-06-19 mgmt shared settings api-status availability(1)

mgmt shared settings api-status log resource-property

NAME

resource-property - Configures the logging of api-status message for the properties that have one of the specified API lifecycle states.

MODULE

mgmt shared settings api-status log

SYNTAX

Configure the resource-property component within the mgmt shared settings api-status log module using the syntax shown in the following sections.

MODIFY

modify resource-property

options:

deprecatedApiAllowed [true | false]

earlyAccessApiAllowed [true | false]

testOnlyApiAllowed [true | false]

DISPLAY

list resource-property

options:

deprecatedApiAllowed

earlyAccessApiAllowed

testOnlyApiAllowed

DESCRIPTION

There are six F5 API lifecycle states namely: internal, testOnly, earlyAccess, generalAccess, deprecated, noStatus. These states are described at the end of this page. For the properties in deprecated, earlyAccess and testOnly states the log::resource-property can be configured by the user.

The resource settings indicate if the user attempt to access the command would generate a [api-status-warning]. These warning will be generated at stderr and /var/log/itm. For example: if the deprecatedApiAllowed property is set to true; then it implies that the accessing the command with deprecated status will generate an [api-status-warning].

Warnings will always be generated regardless of the log::resource-property setting if the property is not available as per the mgmt::shared::settings::api-status::log::availability setting

If these settings are modified in the tmsh interactive mode then, the user has to exit the tmsh interactive mode and re-enter for the settings to be effective.

OPTIONS

deprecatedApiAllowed

Determines if the [api-status-warning] messages will be generated for the properties with deprecated state. The default value is true.

earlyAccessApiAllowed

Determines if the [api-status-warning] messages will be generated for the properties with earlyAccess state. The default value is true.

testOnlyApiAllowed

Determines if the [api-status-warning] messages will be generated for the properties with testOnly state. The default value is true.

The remaining properties are read-only and are not user modifiable.

F5 API lifecycle policy

The F5 API lifecycle consists of five states, namely: internal, testOnly, earlyAccess, generalAccess, deprecated, noStatus. These states can be described as follows:

internal

Internal routing not exposed through public interface

testOnly

Enabled only when test variable configured. Used internally for test cases that require special-case workers to complete functional testing for features not exposed in public API.

earlyAccess

Experimental and susceptible to change in future releases. Used for new features that haven't had time to be solidified.

generalAccess

API that satisfies the API general access release policy.

deprecated

Resource still exists in the API but indicates that there is an preferred alternative or that it may be removed in the future based on deprecation policy.

noStatus

For backward compatibility for APIs not categorized. At some point this status will be disallowed and build time break will occur if resource not categorized with appropriate flag.

In the write-up above the term property is used in place of API to describe the behavior of the mgmt::shared::settings::api-status::log::resource-property settings.

SEE ALSO

mgmt shared settings api-status availability, mgmt shared setting api-status log resource

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2017-06mgmt shared settings api-status log resource-property(1)

mgmt shared settings api-status log resource

NAME

resource - Configures the logging of api-status message for the commands that have one of the specified API lifecycle states.

MODULE

mgmt shared settings api-status log

SYNTAX

Configure the resource component within the mgmt shared settings api-status log module using the syntax shown in the following sections.

MODIFY

modify resource

options:

deprecatedApiAllowed [true | false]

earlyAccessApiAllowed [true | false]

testOnlyApiAllowed [true | false]

DISPLAY

list resource

options:

deprecatedApiAllowed

earlyAccessApiAllowed

testOnlyApiAllowed

DESCRIPTION

There are six F5 API lifecycle states namely: internal, testOnly, earlyAccess, generalAccess, deprecated, noStatus. These states are described at the end of this page. For the commands, command properties in deprecated, earlyAccess and testOnly states the resource can be configured by the user.

The resource settings indicate if the user attempt to access the command would generate a [api-status-warning]. These warning will be generated at stderr and /var/log/itm. For example: if the deprecatedApiAllowed property is set to true; then it implies that the accessing the command with deprecated status will generate an [api-status-warning].

Warnings will always be generated regardless of the log resource setting if the command is not available as per the mgmt::shared::settings::api-status::log::availability setting

If these settings are modified in the tmsh interactive mode then, the user has to exit the tmsh interactive mode and re-enter for the settings to be effective.

OPTIONS

deprecatedApiAllowed

Determines if the [api-status-warning] messages will be generated for the commands with deprecated state. The default value is true.

earlyAccessApiAllowed

Determines if the [api-status-warning] messages will be generated for the commands with earlyAccess state. The default value is true.

testOnlyApiAllowed

Determines if the [api-status-warning] messages will be generated for the commands with testOnly state. The default value is true.

The remaining properties are read-only and are not user modifiable.

F5 API lifecycle policy

The F5 API lifecycle consists of five states, namely: internal, testOnly, earlyAccess, generalAccess, deprecated, noStatus. These states can be described as follows:

internal

Internal routing not exposed through public interface

testOnly

Enabled only when test variable configured. Used internally for test cases that require special-case workers to complete functional testing for features not exposed in public API.

earlyAccess

Experimental and susceptible to change in future releases. Used for new features that haven't had time to be solidified.

generalAccess

API that satisfies the API general access release policy.

deprecated

Resource still exists in the API but indicates that there is an preferred alternative or that it may be removed in the future based on deprecation policy.

noStatus

For backward compatibility for APIs not categorized. At some point this status will be disallowed and build time break will occur if resource not categorized with appropriate flag.

In the write-up above the terms command is used in place of API to describe the behavior of the mgmt::shared::settings::api-status::log::resource settings.

SEE ALSO

mgmt shared settings api-status availability, mgmt shared setting api-status log resource-property

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2017-06-22 mgmt shared settings api-status log resource(1)

net

net address-list

NAME

address-list - An address-list is a list of IP addresses.

MODULE

net

SYNTAX

CREATE/MODIFY

create address-list [name]

modify address-list [[name] | all]

options:

addresses [add | delete | modify | replace-all-with] {
[[ip address]]
}

app-service [name]

description [string]

edit address-list [[name] | all]

options:

all-properties

non-default-properties

DISPLAY

list address-list [[name] | all | [property]]

DELETE
delete address-list [[name] | all]

DESCRIPTION

You can use the address-list component to define reusable lists of addresses.

EXAMPLES

```
create address-list alist1 addresses add { 10.10.1.1 10.10.1.2 192.168.24.0/24 }
```

Creates a new address list, "alist1," with two IPv4 addresses and one IPv4 subnet.

```
modify address-list alist1 addresses modify { 10.10.1.1 { description "management IP at wwmed site3" } }
```

Modifies the above address list with a description for the first address.

```
modify address-list alist1 addresses add { 2001:DB8:a::/64 }
```

Modifies the same address list by adding an IPv6 subnet.

```
list address-list alist1
net address-list alist1 {
  addresses {
    10.10.1.1 {
      description "management IP at wwmed site3"
    }
    10.10.1.2 { }
    192.168.24.0/24 { }
    2001:db8:a::/64 { }
  }
}
```

Shows the modified address list.

OPTIONS

addresses

Specifies a list of individual IP addresses and/or subnets. The format for an IPv4 address is a.b.c.d[/prefix]. The general format for an IPv6 address is a:b:c:d:e:f:g:h[/prefix]; you can shorten this by eliminating leading zeros from each field (for example, you can shorten "2001:0db7:3f4a:09dd:ca90:ff00:0042:8329" to "2001:db7:3f4a:9dd:ca90:ff00:42:8329"), and/or by removing the longest contiguous field of zeros (for example, you can shorten "2001:0:0:0:c34a:0:23ff:678" to "2001::c34a:0:23ff:678"). TMSH accepts any valid text representation of IPv6 addresses, as defined in RFC 2373 (see).

The next keyword specifies the action to take with the addresses (add, delete, modify, or replace the current set of addresses).

add Creates a new address list, which you specify next with IP addresses and/or prefixes in curly braces ({}).

delete

Deletes the address(es) that you specify next, in curly braces ({}).

modify

Makes it possible to replace the optional description(s) for the address(es). You can specify a description in a nested set of curly braces after each address.

replace-all-with

Replaces the current set of IP addresses with the address(es) that you specify next, in curly braces ({}).

app-service

Associates this address list with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers.

description

Is your description for this address list.

SEE ALSO

edit, list, modify, ltm virtual, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2017-2018. All rights reserved.

BIG-IP 2019-04-25 net address-list(1)

net arp

NAME

arp - Manages entries in the Address Resolution Protocol (ARP) table.

MODULE

net

SYNTAX

Configure the arp component within the net module using the syntax in the following sections.

CREATE/MODIFY

create arp [name]

options:

description [string]

ip-address [ip address ... ip address]

mac-address [mac address]

modify arp [name]

options:

description [string]

mac-address [mac address]

edit arp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list arp

list arp [[[name] | [glob] | [regex]] ...]

show running-config arp

show running-config arp

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show arp

show arp [[[name] | [glob] | [regex]] ...]

options:

(dynamic | static)

DELETE

delete arp [name]

DESCRIPTION

You can use the arp component to add entries to or delete entries from the ARP table.

You can create static ARP entries for IPv4 addresses to link-layer addresses, such as Ethernet media access control (MAC) addresses. You can view and delete static and dynamic ARP entries.

Note that you can use the db component in the sys module to configure how the system handles ARP entries for dynamic timeout, maximum dynamic entries, add reciprocal, and maximum retries. For more information, see sys db.

EXAMPLES

```
create arp myARP ip-address 10.10.10.20 mac-address 00:0b:09:88:00:9a
```

Creates an arp mapping of the IP address 10.10.10.20 to the MAC address 00:0b:09:88:00:9a, and the name of this entry is myARP. Alternatively, the address can be used as the name, like the following example.

```
create arp 10.10.10.20 mac-address 00:0b:09::88:00:9a
```

Creates an arp mapping of IP address 10.10.10.20 to the MAC address 00:0b:09:88:00:9a.

```
modify arp 10.10.10.20 mac-address 00:0b:09:88:00:9b
```

Modifies the ARP mapping of the ARP entry named 10.10.10.20 to the MAC address 00:0b:09:88:00:9b.

```
show arp
```

Displays ARP status and statistics for the system.

```
show arp any%2
```

Displays ARP status and statistics for all IP addresses in route domain 2. A glob expression displays the same result: show arp *%2.

```
list arp all-properties
```

Displays all properties for all ARP entries for the system.

```
list arp non-default-properties
```

Displays all non-default properties for all ARP entries for the system.

delete arp all

Deletes all ARP entries for the system.

delete arp myARP

Deletes the ARP entry named myARP.

OPTIONS

description
User defined description.

dynamic
Displays the status of dynamic ARP entries.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ip address
Specifies the IP address, in one of the following formats, for which you want to configure an ARP entry:

IPv4 address in dotted-quad notation, for example, 10.10.10.1
host name, for example, www.f5.com

You can also specify a list of IP addresses separated by a single space. For example, this list contains three IP addresses: 10.10.10.20 10.10.10.21 10.10.10.22.

ip-address
The IP address to be mapped. This is optional, and if not present, the name needs to be a string that represents a valid IP address.

mac-address
Specifies a 6-byte ethernet address in not case-sensitive hexadecimal colon notation, for example, 00:0b:09:88:00:9a. You must specify a MAC address when you create an ARP entry.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

static
Displays the status of static ARP entries.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2016-03-24 net arp(1)

net bwc policy

NAME

policy - Configures a bandwidth control policy for traffic flow.

MODULE

net bwc

SYNTAX

Configure the policy component within the net bwc module using the syntax in the following sections.

CREATE/MODIFY

create policy [name]
modify policy [name]

options:

app-service [[string] | none]
description [string]
dynamic [enabled]
max-rate [integer]
max-user-rate [integer]
max-user-rate-pps [integer]
ip-tos [integer | pass-through]

```

link-qos [integer | pass-through]
measure [ disabled ]
log-publisher [[string] | none]
log-period [integer]
categories [none] {
  max-cat-rate [integer]
  max-cat-rate-percentage [integer]
  ip-tos [ integer | pass-through]
  link-qos [integer | pass-through]
  traffic-priority-map [string]
}
traffic-priority-map [string]

```

```

edit policy [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

DISPLAY

```

list policy
list policy [ [ [name] | [glob] | [regex] ] ... ]
show running-config net policy
show running-config net policy [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line

```

DELETE

```

delete net policy [all | [name] ]

```

DESCRIPTION

You can use the net bwc policy to create a bandwidth control policy to handle traffic flow, and then associate it with other components such as packet filter, iRule and virtual server. For details on packet filter, virtual server, please refer to the respective documentation.

EXAMPLES

```

create net bwc policy

```

Creates a bwc policy (see below).

```

list net bwc policy all-properties

```

Displays all of the properties of all of the bwc policies.

```

delete net bwc policy

```

Deletes a policy (see below).

Example for static policy:

```

net bwc policy silver_static_policy {
  max-rate 120mbps
}

```

Example for dynamic policy:

```

net bwc policy gold-dynamic-policy {
  categories {
    web {
      description "This is a web test category."
      max-cat-rate 600kbps
      ip-tos 7
      link-qos 5
    }
  }
  description "This is a test."
  dynamic enabled
  max-rate 40gbps
  max-user-rate 1gbps
}

```

Example for dynamic policy with measure enabled:

```

net bwc policy gold-dynamic-policy {
  categories {
    web {
      description "This is a web test category."
      max-cat-rate 600kbps
      ip-tos 7
      link-qos 5
    }
  }
  description "This is a test."
  dynamic enabled
  measure enabled
  log_publisher /Common/my_log_publisher
}

```

```

log-period 2048
max-rate 40gbps
max-user-rate 1gbps
}

```

Example for BWC using packet filter:

```

net bwc policy bwc {
  max-rate 1mbps
}

```

Define packet filter with bwc on it:

```

net packet-filter pfilter {
  action continue
  bwc policy bwc
  logging enabled
  order 2
  rule ip
}

```

Example for BWC association with virtual server:

```

itm virtual l2-for-virtual {
  destination 0.0.0.0:any
  l2-forward
  mask any
  profiles {
    fastL4 { }
  }
  rules {
    bwc_test
  }
  translate-address disabled
  translate-port disabled
  vlans {
    lan
    wan
  }
  vlans-enabled
}
itm virtual tcp-passthrough {
  destination 0.0.0.0:http
  ip-protocol tcp
  mask any
  profiles {
    tcp { }
  }
  rules {
    bwc_test
  }
  translate-address disabled
  vlans-disabled
}

```

Example for Delete bwc policy:

```

net bwc policy silver_static_policy

```

Example for bwc policy traffic map:

```

net bwc policy bwc-policy-105 {
  categories {
    cat1 {
      max-cat-rate 10mbps
      traffic-priority-map tc1->cat1
    }
    cat2 {
      max-cat-rate 10mbps
      traffic-priority-map tc1->cat2
    }
  }
  dynamic enabled
  max-rate 100mbps
  max-user-rate 10mbps
}

```

Notes: Only static policies are supported for association with packet filter or virtual server components.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

name Specifies a unique name for the policy. This option is required for the commands create, delete, and modify.

description

User defined description.

max-rate

Specifies the maximum bandwidth that traffic is allowed using the policy associated. The range is from 1Mbps to 320Gbps.

Valid units: bps(default), gbps, kbps, mbps.

max-user-rate

Specifies the maximum bandwidth that traffic is allowed using the policy associated. The range is from 5kbps to 2Gbps.

Valid units: bps(default), gbps, kbps, mbps.

max-user-rate-pps

Specifies the limiter in packets per second that traffic is allowed using the associated policy. This does not allocate any fairshare bandwidth. When configured this acts purely as a simple packet limiter. It is packet size and protocol agnostic. It can be configured only on a dynamic policy. When configured along with mbps values, whichever lower limit pps vs mbps is applied. When configured, both need to pass for packets to go through. The default value is 0 (not configured).

traffic-priority-map

Specifies the bwc priority-group to use during congestion. This is optional and to be configured only as needed. A bwc priority-group can be shared and amongst categories of the same bwc policy but not across bwc policies. When configured on policy or category, the max-user-rate or max-cat-rate as configured would be additionally applied. Thus the lower of all values for max. rates would take effect.

max-cat-rate

Specifies the maximum bandwidth that traffic is allowed using this category with associated policy. The range is from 5Kbps to max-user-rate.

Valid units: bps(default), gbps, kbps, mbps.

max-cat-rate-percentage

Specifies the percentage of the value of the max-cat-rate option of the category, which is associated with the net bwc policy component to which this shaping policy is associated, that is available for this traffic flow. It is the maximum bandwidth as percentage of that traffic is allowed using this category with associated policy. The range is from 1 to 100.

dynamic

Specifies the type for policy to be dynamic type. This option is optional for the commands create, delete, and modify. The default valid is disabled. When dynamic is disabled, the policy type is said to be static, where the maximum rate is enforced for combined traffic using the policy and no fairness bandwidth guarantee for each of the traffic respectively. The default value is: disabled. Note: policy type change modification is a disallowed configuration.

By enabling this option, the policy is dynamic type and requires you to configure max-user-rate-range. This type of policy enforces fairness for all the traffic associated with the policy and also for each traffic within the policy.

ip-tos

Specifies an IP ToS number for the traffic using the net bwc policy. This option specifies the ToS level that the traffic management system assigns to UDP packets when sending them. The default value is pass-through, which indicates, do not modify UDP packets. The valid range for IP ToS value that can be specified is 0 to 63.

Note: If this is specified, bandwidth policy is not enforced. The packets are just marked for a downstream system to process.

link-qos

Specifies a Link QoS (VLAN priority) for the traffic using the net bwc policy. This option specifies the QoS level that the system assigns to UDP packets when sending. The default value is pass-through, which indicates, do not modify UDP packets. The valid range for QoS value is 0 to 7.

Note: If this is specified, bandwidth policy is not enforced. The packets are just marked for a downstream system to process.

measure

Enables or disables bandwidth measurement on all the future instances of bwc policy. Users can override this setting using iRules. If enabling measurement on all instances is not desired then users can keep this setting disabled and use iRules to enable measurement on specific instances of bwc policy.

log_publisher

Specifies the name of the log publisher configured in the system. Bandwidth measurement results will be sent to this log publisher.

log_period

Time interval in milliseconds representing the frequency of generation of bandwidth measurement logs.

categories

This specifies the categories under policy. Note: policy need to be enabled as dynamic to configure categories. Up to a maximum of 32 categories can be configured. All the categories under the dynamic policy share the bandwidth as specified for the category, up to a maximum of max-user-rate. Specify the maximum bandwidth for the category of traffic using max-cat-rate or by max-cat-rate-percentage as a percentage of the maximum user rate. Either only the range or absolute value is required.

Example to configure a dynamic bandwidth policy category using tmsh:

```
root@(localhost)(cfg-sync
Standalone)(Active)/(Common)(tmos.net.bwc.policy.gold-dynamic-policy)# categories add { web { max-cat-rate 600kbps } }

net bwc policy gold-dynamic-policy {
categories {
  web {
max-cat-rate 600kbps
  }
}
dynamic enabled
max-rate 40gbps
max-user-rate 1gbps
}
```

The parameters for dynamic policy and categories:

```
net bwc policy test-policy {
app-service none
categories {
  web {
app-service none
description "This is a web test cat"
max-cat-rate 600kbps
max-cat-rate-percentage 0
ip-tos 6
  }
}
description "This is a test"
dynamic enabled
ip-tos pass-through
link-qos pass-through
max-rate 40gbps
max-user-rate 1gbps
measure enabled
log-publisher /Common/my_log_publisher
log-period 2048
partition Common
}
```

Few Examples using iRule:

Please refer to iRule documentation for complete list of bwc commands. Below are few examples and do not cover all cases.

Example to associate static bwc policy using iRule:

```
when CLIENT_ACCEPTED {
  BWC::policy attach silver_static_policy
}
```

Example to associate dynamic bwc policy using iRule:

```
when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]
  BWC::policy attach gold-dynamic-policy $mycookie
}
```

Example for bwc policy to mark traffic flows using iRule:

```
BWC::mark > >
```

So to assign a policy, color, and mark here is an example rule

```
when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]:[TCP::remote_port]
  BWC::policy attach gold_user $mycookie
  BWC::color set gold_user p2p
  BWC::mark set gold_user tos 8 qos 4
}
```

Example for using bwc policy category to color a flow using iRule:

After a flow has been assigned a policy, at some later time when the traffic is classified the user can assign an application to this flow. This uses the bwc config to create a bwc policy with the categories keyword: for example, p2p category below:

```
tmsh create net bwc policy gold_user categories add { p2p { max-cat-rate 8mbps } } max-rate 10mbps max-user-rate 10mbps dynamic enabled
```

The rule args

```
BWC::color
```

So to assign a policy and color here is an example rule

```

when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]:[TCP::remote_port]
  BWC::policy attach gold_user $mycookie
  BWC::color set gold_user p2p
}

```

Example for bwc policy rate change using iRule:

After a policy is created, irule can modify the rate for a session or category

The rule args

```
BWC::rate
```

```
BWC::rate
```

So to modify the rate

```

when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]:[TCP::remote_port]
  BWC::policy attach gold_user $mycookie
  BWC::color set gold_user p2p
  BWC::mark set gold_user tos 8 qos 4
  BWC::rate $mycookie p2p 1000000bps
}

```

Example for bwc policy to measure the bandwidth using iRule:

```

BWC::measure << | | >
[session_str]>

```

To start the bandwidth measurement for BWC policy

```

when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]:[TCP::remote_port]
  BWC::policy attach gold_user $mycookie
  BWC::measure start session
}

```

The above iRule will start the measurement of bandwidth on the gold_user policy instance. The results will be published to the destination specified in the log_publisher setting for the gold_user bwc policy. The measurement results will be logged every 'log_period' amount of time, which is also specified in the policy settings for gold_user.

Note: Attaching a BWC policy is a pre-requisite for all 'BWC::measure.' iRules. Failing to do so will result in the iRule execution failure which in turn will abort the connection.

To start the bandwidth measurement for a flow inside a BWC policy and tag the results.

```

when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]:[TCP::remote_port]
  BWC::policy attach gold_user $mycookie
  BWC::measure identifier MYFLOW flow
  BWC::measure start flow
}

```

The above iRule will start the measurement of bandwidth on the current flow or the flow over which the current iRule is running. The results will be published to the destination specified in the log_publisher setting for the gold_user bwc policy. The measurement results will be logged every 'log_period' amount of time, which is also specified in the policy settings for gold_user. Every log message containing the bandwidth result will carry the tag 'MYFLOW'. This helps identify different types of bandwidth results when analyzing the bandwidth measurement results.

To start the bandwidth measurement for a BWC session and get the periodic results.

```

when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]:[TCP::remote_port]
  set measureID "MYSESSION"
  BWC::policy attach gold_user $mycookie
  BWC::measure identifier $measureID session
  BWC::measure start session
}

```

```

when SERVER_CONNECTED {
  TCP::collect
  set count 0
}

```

```

when SERVER_DATA {
  if {$count >= 1000 } {
    set rate [BWC::measure get rate session]
    set bytes [BWC::measure get bytes session]
  }
}

```

```

log local0. "Rate $rate/sec : Bytes $bytes : for address $mycookie"
set count 0
}
TCP::release

```

```
TCP::collect
incr count
}
```

The above example creates an instance of BWC policy `gold_user` and enables measurement on it upon the iRule event `CLIENT_ACCEPTED`. It also starts collecting so that it keeps getting notified upon the data arrival. For every 1000 packets a measurement of bandwidth is logged. This is an example that illustrates how to measure bandwidth periodically using iRules. Note that the bandwidth measurement results are still sent to configured `log_publisher` at every `log_period` interval.

An example log message that is sent to a log publisher.

```
Apr 2 16:29:04 MYSESSION BWC Measurement: Moving average - 539277 bytes/sec. Total bytes - 7305051
```

Example for `bwc` policy using `pps`:

```
BWC::pps

when CLIENT_ACCEPTED {
  set mycookie [IP::remote_addr]:[TCP::remote_port]
  BWC::policy attach gold_user $mycookie
  BWC::pps 100
}
```

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2016. All rights reserved.

BIG-IP 2017-05-18 net bwc policy(1)

net bwc priority-group

NAME

`priority-group` - Configures a bandwidth control policy for traffic flow.

MODULE

`net bwc`

SYNTAX

Configure the `priority-group` component within the `net bwc` module using the syntax in the following sections.

CREATE/MODIFY

```
create priority-group [name]
modify priority-group [name]
options:
  app-service [[string] | none]
  description [string]
  priority-classes [none] {
    description [string]
    weight-percentage [integer]
  }
}
```

```
edit priority-group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list priority-group
list priority-group [ [ [name] | [glob] | [regex] ] ... ]
show running-config net bwc priority-group
show running-config net bwc priority-group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

DELETE

```
delete net bwc priority-group [all | [name] ]
```

DESCRIPTION

You can use the `net bwc priority-group` to create a bandwidth traffic group policy to handle traffic flow, and then associate it with `bwc` component such as `bwc policy/category`.

EXAMPLES

```
create net bwc priority-group
```

Creates a bwc policy (see below).

```
list net bwc priority-group all-properties
```

Displays all of the properties of all of the bwc policies.

```
delete net bwc priority-group
```

Deletes a bwc traffic group (see below).

Example for dynamic policy:

```
net bwc priority-group tc-GOLD {
  priority-classes {
    tc-BLUE {
      weight-percentage 30
    }
    tc-RED {
      weight-percentage 20
    }
  }
}
```

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

name Specifies a unique name for the policy. This option is required for the commands create, delete, and modify.

description

User defined description.

weight-percentage

Specifies the percentage of maximum bandwidth that traffic is allowed during congestion using the traffic class associated. The valid range is 5-100.

iRule

Please refer to iRule documentation for iRule to use bandwidth control policy.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2014. All rights reserved.

BIG-IP 2016-05-18 net bwc priority-group(1)

net bwc traffic-group

NAME

traffic-group - Configures a bandwidth control policy for traffic flow.

MODULE

net bwc

SYNTAX

Configure the traffic-group component within the net bwc module using the syntax in the following sections.

CREATE/MODIFY

```
create traffic-group [name]
```

```
modify traffic-group [name]
```

options:

```
app-service [[string] | none]
```

```
description [string]
```

```
dynamic [ enabled ]
```

```
priority-classes [none] {
```

```
  weight-percentage [integer]
```

```
}
```

```
edit traffic-group [ [ [name] | [glob] | [regex] ] ... ]
```

options:

- all-properties
- non-default-properties

DISPLAY

```
list traffic-group
```

```
list traffic-group [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config net bwc traffic-group
```

```
show running-config net bwc traffic-group [ [ [name] | [glob] | [regex] ] ... ]
```

options:

- all-properties
- non-default-properties
- one-line

DELETE

```
delete net bwc traffic-group [all | [name] ]
```

DESCRIPTION

You can use the net bwc traffic-group to create a bandwidth traffic group policy to handle traffic flow, and then associate it with bwc component such as bwc policy/category.

EXAMPLES

```
create net bwc traffic-group
```

Creates a bwc policy (see below).

```
list net bwc traffic-group all-properties
```

Displays all of the properties of all of the bwc policies.

```
delete net bwc traffic-group
```

Deletes a bwc traffic group (see below).

Example for dynamic policy:

```
net bwc traffic-group tc-GOLD {
  priority-classes {
    tc-BLUE {
      weight-percentage 30
    }
    tc-RED {
      weight-percentage 20
    }
  }
}
```

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

name Specifies a unique name for the policy. This option is required for the commands create, delete, and modify.

description

User defined description.

weight-percentage

Specifies the maximum bandwidth that traffic is allowed during congestion using the traffic class associated. The range is from 1Mbps to 320Gbps.

Valid units: bps(default), gbps, kbps, mbps.

dynamic

Specifies the type for traffic group to be dynamic type. This option is optional for the commands create, delete, and modify. The default valid is disabled. When dynamic is disabled, the traffic group type is said to be static. The type of traffic group should match with be bwc policy that would be mapped. Note: policy type change modification is a disallowed configuration.

iRule

Please refer to iRule documentation for iRule to use bandwidth control policy.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015. All rights reserved.

net clone-stats

NAME

clone-stats - Display and Reset Clone Statistics.

MODULE

net

SYNTAX

Display and Reset the clone-stats component within the net module using the syntax in the following section.

MODIFY

reset-stats clone-stats

DISPLAY

show clone-stats

options:

(default | field-fmt)

DESCRIPTION

You can use the clone-stats component to display and reset clone statistics. This will help to debug the clone feature.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013, 2016. All rights reserved.

BIG-IP 2017-01-20 net clone-stats(1)

net cmetrics

NAME

cmetrics - Displays and deletes entries in the route metrics table on the BIG-IP(r) system.

MODULE

net

SYNTAX

Use the cmetrics component within the net module to view route metrics or delete a route metric entry using the following syntax.

DISPLAY

show cmetrics

option:

bandwidth

dest-addr [ip address]

hwaddress

mtu

rtt

rttvar

ssthresh

tmm

DELETE

delete cmetrics

option:

dest-addr [IP address]

DESCRIPTION

You can use the cmetrics component to display entries in the route metrics table on the BIG-IP system. Additionally, you can delete a specified route metric entry from the table. The options are display-only values and cannot be used for filtering.

Note: You can delete only entries that have no connection references.

EXAMPLES

```
show cmetrics
```

Displays all the entries in the route metrics table.

```
delete cmetrics dest-addr 10.10.1.11
```

Deletes the entry with destination IP address 10.10.1.11 from the route metrics table.

OPTIONS

bandwidth

Displays the size of the channel in kbps. Computed as cwnd/rtt.

dest-addr

Specifies the destination IP address of the entry that you want to display or delete. You can enter this address in either IPv4 or IPv6 format.

hwaddress

Displays the Media Access Control (MAC) address for the route.

mtu Displays the maximum transmit unit size (in bytes) on the route.

rtt Displays the round-trip time on the route in units of 100ns.

rttvar

Displays the variation in the round-trip time in units of 100ns.

ssthresh

Displays the cached slow-start threshold in bytes.

tmm Displays the identifying number of the tmm (Traffic Management Microkernel).

SEE ALSO

delete, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2015-10-08 net cmetrics(1)

net cos global-settings

NAME

global-settings - Configures the global configuration for class of service (CoS).

MODULE

net cos

SYNTAX

Modify the global-settings component within the net cos module using the syntax shown in the following sections.

CREATE/MODIFY

```
modify global-settings
options:
  feature-enabled
  feature-disabled
  precedence [dscp-only, 8021p-only]
  default-map-dscp
  [add | delete | modify | replace-all-with] {
  [map-dscp-name] ...
  }
  default-map-8021p
  [add | delete | modify | replace-all-with] {
  [map-8021p-name] ...
  }
  default-traffic-priority [ traffic-priority-name ]
```

DISPLAY

```
list global-settings
options:
  all-properties
  non-default-properties
```

one-line
show global-settings

DESCRIPTION

You can use the global-settings component to configure and view information about the global settings of all CoS behavior.

show keyword displays an analysis of the relative weights of the associated traffic-priority objects.

EXAMPLES

```
modify global-settings default-traffic-priority NORMAL_PRIORITY
```

Replace the default traffic-priority with traffic-priority NORMAL_PRIORITY.

```
modify global-settings default-map-8021p add { VOIP }
```

Add the VOIP 802.1p mapping. The VOIP object specifies the 802.1p field value and associated traffic priority.

OPTIONS

feature-enabled

Enable 8 hardware egress CoS queue feature.

feature-disabled

Disable 8 hardware egress CoS queue feature.

precedence

Specifies the precedence between handling of DSCP and 802.1p. Currently, provided options are dscp-only and 8021p-only.

default-map-dscp

Enables adding and removal of mappings between DSCP field values and traffic priorities. See net cos traffic-priority and net cos map-dscp.

default-map-8021p

Enables adding and removal of mappings between 802.1p field values and traffic priorities.

default-traffic-priority

Specifies the default traffic-priority which is applied to all traffic that does not match a specified DSCP/802.1p field value. This allows the user to specify only the mappings which do not match the default.

SEE ALSO

net cos traffic-priority, net cos map-dscp

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 net cos global-settings(1)

net cos map-8021p

NAME

map-8021p - Configures vlan 8021.p tag to traffic priority mapping.

MODULE

net cos

SYNTAX

Modify the map-8021p component within the net cos module using the syntax shown in the following sections.

CREATE/MODIFY

```
create map-8021p [name]
```

```
modify map-8021p [name]
```

options:

```
value [0..7]
```

```
traffic-priority [name]
```

```
edit map-8021p [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list map-8021p
```

```
list map-8021p [ [ [name] | [glob] | [regex] ] ... ]
```

```
show map-8021p
show map-8021p [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

```
DELETE
delete map-8021p [name]
```

DESCRIPTION

The map-8021p object allows users to associate 802.1p field values to relative traffic priority. These objects are associated with active system configuration via net cos global-settings.

EXAMPLES

```
create map-8021p VOIP value 4 traffic-priority HIGH_PRIORITY
```

Create the map-8021p named VOIP that associates 802.1p value 4 traffic with traffic-priority named HIGH_PRIORITY.

```
delete map-8021p VOIP
```

Delete the map-8021p named VOIP.

OPTIONS

```
value
Specifies the 802.1p field value.
```

```
traffic-priority
Specifies the traffic-priority object associated with traffic matching value.
```

SEE ALSO

create, delete, edit, glob, list, net cos global-settings, modify, net cos traffic-priority, net cos map-dscp, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-07-01 net cos map-8021p(1)

net cos map-dscp

NAME

map-dscp - Configures IP DSCP field to traffic priority mapping.

MODULE

net cos

SYNTAX

Modify the map-dscp component within the net cos module using the syntax shown in the following sections.

```
CREATE/MODIFY
create map-dscp [name]
modify map-dscp [name]
options:
  value [0..7]
  traffic-priority [name]
```

```
edit map-dscp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list map-dscp
list map-dscp [ [ [name] | [glob] | [regex] ] ... ]
```

```
show map-dscp
show map-dscp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DELETE
delete map-dscp [name]

DESCRIPTION

The map-dscp object allows users to associate DSCP field values to relative traffic priority. These objects are associated with active system configuration via net cos global-settings.

EXAMPLES

```
create map-dscp VOIP value 4 traffic-priority HIGH_PRIORITY
```

Create the map-dscp named VOIP that associates DSCP value 4 traffic with traffic-priority named HIGH_PRIORITY.

```
delete map-dscp VOIP
```

Delete the map-dscp named VOIP.

OPTIONS

value
Specifies the DSCP field value.

traffic-priority
Specifies the traffic-priority object associated with traffic matching value.

SEE ALSO

create, delete, edit, glob, list, net cos global-settings, modify, net cos traffic-priority, net cos map-dscp, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-07-01 net cos map-dscp(1)

net cos traffic-priority

NAME

traffic-priority - Configures a traffic priority object.

MODULE

net cos

SYNTAX

Modify the traffic-priority component within the net cos module using the syntax shown in the following sections.

CREATE/MODIFY

```
create traffic-priority [name]
modify traffic-priority [name]
options:
  weight [1..127]
  buffer [1..127]
```

```
edit traffic-priority [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list traffic-priority
list traffic-priority [ [ [name] | [glob] | [regex] ] ... ]

show traffic-priority
show traffic-priority [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DELETE

```
delete traffic-priority [name]
```

DESCRIPTION

The traffic-priority object allows users to assign relative scheduling and buffer weightings. These objects are associated to specific traffic with net cos map-dscp and net cos map-8021p. There can be at most 8 traffic-priorities defined in the system. The DEFAULT_PRIORITY priority may be deleted or modified as desired.

EXAMPLES

```
create traffic-priority HIGH_PRIORITY weight 127
```

Create the traffic-priority HIGH_PRIORITY that has a weight of 127.

```
delete traffic-priority HIGH_PRIORITY
```

Delete the traffic-priority named HIGH_PRIORITY.

OPTIONS

weight

Specifies the egress buffer weight. This value is used relative to other egress traffic-priority objects typical of weighted round-robin behavior.

buffer

Specifies the relative buffer weight where available egress buffer space is distributed with consistent relative weight.

SEE ALSO

create, delete, edit, glob, list, net cos global-settings, modify, net cos map-dscp, net cos map-8021p, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 net cos traffic-priority(1)

net dag-globals

NAME

dag-globals - configure global disaggregation settings.

MODULE

net

SYNTAX

Configure the dag-globals component within the net module using the syntax shown in the following sections.

MODIFY

```
modify dag-globals
```

options:

```
round-robin-mode [global | local]
```

```
dag-ipv6-prefix-len [integer]
```

```
icmp-hash [icmp | ipicmp]
```

```
icmp-monitor-priority [high | normal]
```

```
edit dag-globals
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list dag-globals
```

```
show running-config dag-globals
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DESCRIPTION

Provides the ability to configure global disaggregation settings.

round-robin-mode

Specifies whether the round robin disaggregator (DAG) on a blade can disaggregate packets to all the TMMs in the system or only to the TMMs local to the blade.

dag-ipv6-prefix-len

Specifies whether SPDAG or IPv6 prefix DAG should be used to disaggregate IPv6 traffic when vlan cmp hash is set to src-ip or dst-ip. The default value is 128, using SPDAG.

icmp-hash

Specifies ICMP hash for ICMP echo request and ICMP echo reply in SW DAG. ICMP echo request and ICMP echo reply can be disaggregated based on ICMP id or ICMP id and IP addresses. The default ICMP hash is ICMP id.

icmp-monitor-priority
Specifies the priority for ICMP monitor traffic. The default value is normal.

SEE ALSO

list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2015. All rights reserved.

BIG-IP 2018-05-09 net dag-globals(1)

net dns-resolver

NAME

resolver - Configures a DNS resolver on the BIG-IP(r) system.

MODULE

net dns-resolver

SYNTAX

Configure the DNS resolver component using the syntax in the following sections.

CREATE/MODIFY

```
create [name]
modify [name]
options:
  answer-default-zones [yes | no]
  app-service [[string] | none]
  cache-size [integer]
  description [string]
  forward-zones [add | delete | modify | replace-all-with] {
    [ [zone-name] ] {
options:
  nameservers [add | delete | replace-all-with] {
    [ [IPv4address:port] | [IPv6address.port] ]
  }
  nameservers none
  }
  forward-zones none
  randomize-query-name-case [yes | no]
  route-domain [name]
  use-ipv4 [yes | no]
  use-ipv6 [yes | no]
  use-tcp [yes | no]
  use-udp [yes | no]
```

DISPLAY

```
list
list [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
show [name]
reset-stats
```

DELETE

```
delete [name]
```

DESCRIPTION

You can use the dns-resolver component to configure and view information about a DNS Resolver object. A DNS resolver performs recursive resolution to fill its cache.

Important: When sizing caches, consider the total amount of memory available and how you wish to allocate memory for DNS caching. Note that cache sizing values are per-TMM process; therefore, a platform with eight TMMs consumes the amount of memory set for the Resolver object times eight.

Important: DNS Resolver objects use the DNS root nameservers published by InterNIC.

EXAMPLES

```
list
```

Displays the properties of the DNS Resolver myRes.

OPTIONS

answer-default-zones

Specifies whether the resolver answers queries for default zones: localhost, reverse 127.0.0.1 and ::1, and AS112 zones. The default value is no.

app-service

Specifies the name of the application service to which this dns-resolver belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this dns-resolver. Only the application service can modify or delete this dns-resolver.

cache-size

Specifies the maximum cache size in bytes of the DNS Resolver object. The default value is 5767168.

The BIG-IP system caches the supporting records in a DNS response in the resource record cache. After the maximum size of the cache is reached, when new or refreshed content is added to the cache, the expired and older content is removed from the cache. A higher maximum size allows more DNS responses to be cached and increases the cache hit percentage. A lower maximum size forces earlier eviction of cached content, but can lower the cache hit percentage.

description

User defined description.

forward-zones

Adds, deletes, modifies, or replaces a set of forward zones on a DNS Resolver, by specifying zone name(s). A given zone name should only use the symbols allowed for a fully qualified domain name (FQDN), namely ASCII letters a through z, digits 0 through 9, hyphen -, and period .. For example site.example.com would be a valid zone name.

A DNS Resolver configured with a forward zone will forward any queries that resulted in a cache-miss (the answer was not available in the cache) and which also match a configured zone name, to the nameserver specified on the zone. If no nameservers are specified on the zone, an automatic SERVFAIL is returned. When a forward zone's nameserver returns a valid response to the DNS Cache, that response is cached and then returned to the requestor.

nameservers

Adds, deletes, modifies, or replaces a set of nameservers in a forward zone on a DNS Resolver. A nameserver is represented by an IP address and port in the format [IPv4:port] or [IPv6:port], for example 10.10.10.10:53 or 2001::1:ff:53, respectively.

If more than one nameserver is listed for a given forward zone, a matching query will be sent to the nameserver that is currently deemed the most responsive (based on RTTs). If no response is received within a certain window of time, the DNS Resolver will resend the query to another nameserver with an increased wait window, until a response is received.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

randomize-query-name-case

Specifies whether the resolver randomizes the case of query names. The default value is yes.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

route-domain

Specifies the route domain the resolver uses for outbound traffic. The default value is the default route domain.

use-ipv4

Specifies whether the resolver sends DNS queries to IPv4 addresses. The default value is yes.

use-ipv6

Specifies whether the resolver sends DNS queries to IPv6 addresses. The default value is yes.

use-tcp

Specifies whether the resolver can send queries over the TCP protocol. The default value is yes.

use-udp

Specifies whether the resolver can send queries over the UDP protocol. The default value is yes.

SEE ALSO

create, delete, edit, glob, list, show, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

net f5optics

NAME

f5optics - Manage f5optics package.

MODULE

net

SYNTAX

Manage net f5optics using the syntax in the following section.

INSTALL

install f5optics

options:

recovery

slot [[slot number] | all]

SHOW

show f5optics

DESCRIPTION

The f5optics package specifies the supported optical PHY modules and their tuning parameters for network interfaces. You can use the f5optics component to manage the f5optics package. The install command will install f5optics with the latest f5optics package under /shared/f5optics/images. Additionally, for a cluster, you can use the slot option to specify either a specific slot or all slots to install the f5optics package.

Notice that the installation will cause network interfaces to be re-initialized on the affected slot(s), hence a temporary traffic interruption.

The show command will display the version information of the installed f5optics package on the appliance, or all the slots in a cluster.

EXAMPLES

install net f5optics

Install the f5optics package on the appliance, or all the slots in a cluster.

install net f5optics slot all

Install the f5optics package on all the slots in a cluster.

install net f5optics recovery slot all

Install the f5optics package on all the slots in a cluster, even if it is not newer than the installed package.

install net f5optics slot 2

Install the f5optics package on slot 2.

show net f5optics

Show the f5optics package on the appliance, or all the slots in a cluster.

OPTIONS

recovery

Always install the latest f5optics package, even if it is not newer than the installed package. This is for recovery purposes.

slot [[slot number] | all]

Install the f5optics package on either a specific slot or all slots in a cluster, without changing the active volume of the slot(s).

This option is only available in a clustered environment.

SEE ALSO

install, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2016-06-01 net f5optics(1)

net fdb tunnel

NAME

tunnel - Manages tunnel entries in the Layer 2 Forwarding table.

MODULE

net fdb

SYNTAX

Configure the tunnel component within the net fdb module using the syntax in the following sections.

MODIFY

modify tunnel [tunnel name]

options:

app-service [[string] | none]

records [add | delete | modify | replace-all-with] {
[MAC address] {

app-service [[string] | none]

description [string]

endpoint [IP address]

endpoints [add | delete | modify | replace-all-with] {
[IP address(es)]

}

replicators [add | delete | modify | replace-all-with] {
[IP address(es)]

}

}

}

records none

DISPLAY

list tunnel

list tunnel [[[tunnel name] | [glob] | [regex]] ...]

show tunnel

show tunnel [[[tunnel name] | [glob] | [regex]] ...]

options:

all-records

dynamic

field-fmt

static

DELETE

delete tunnel [tunnel name]

options:

all-records

dynamic

static

DESCRIPTION

You can use the tunnel component to manage tunnel entries in the Layer 2 Forwarding table.

EXAMPLES

```
modify tunnel t1 records add { 00:01:02:03:04:05 { endpoint 10.10.0.2 } }
```

Creates a tunnel entry for the specified tunnel in the Layer 2 Forwarding table.

```
modify tunnel t1 records none
```

Deletes all configured tunnel entries for the specified tunnel in the Layer 2 Forwarding table.

```
modify tunnel t1 records add { ff:ff:ff:ff:ff:ff { endpoints add { 10.10.0.2 10.10.0.3 10.10.0.4 } } }
```

Creates a set of endpoints for the specified tunnel in the Layer 2 Forwarding table.

```
list tunnel
```

Displays all configured tunnel entries in the Layer 2 Forwarding table.

```
show tunnel
```

Displays all dynamic and static tunnel entries in the Layer 2 Forwarding table.

```
delete tunnel t1 all-records
```

Deletes all dynamic and static tunnel entries in the Layer 2 Forwarding table.

OPTIONS

all-records

Applies the command to all dynamic and static tunnel entries in the Layer 2 Forwarding table.

dynamic

Applies the command to all dynamic tunnel entries in the Layer 2 Forwarding table.

static

Applies the command to all static tunnel entries in the Layer 2 Forwarding table.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

glob, net tunnels, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2017. All rights reserved.

BIG-IP 2017-05-11 net fdb tunnel(1)

net fdb vlan

NAME

vlan - Manages VLAN entries in the Layer 2 Forwarding table.

MODULE

net fdb

SYNTAX

Configure the vlan component within the net fdb module using the syntax in the following sections.

MODIFY

modify vlan [vlan name]

options:

app-service [[string] | none]

records

[add | delete | modify | replace-all-with] {

[MAC address] ... {

app-service [[string] | none]

description [string]

trunk [trunk name]

interface [interface name]

}

}

records none

edit vlan [[all | [vlan name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list vlan

list vlan [[[vlan name] | [glob] | [regex]] ...]

show running-config vlan

show running-config vlan [[vlan name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

records

show vlan

show vlan [[[vlan name] | [glob] | [regex]] ...]

options:

dynamic

field-fmt

static

DELETE

delete vlan

delete vlan [all | [vlan name]]

options:

all-records

dynamic

static

DESCRIPTION

You can use the vlan component to manage entries in VLAN Layer 2 Forwarding tables.

EXAMPLES

```
modify vlan internal records add { 00:0b:09:88:00:9a { interface 1.2 } }
```

Creates a mapping of the MAC address 00:0b:09:88:00:9a to interface 1.2 on VLAN internal.

```
modify vlan internal records modify { 00:0b:09:88:00:9a { interface 1.1 } }
```

Modifies the mapping of the MAC address 00:0b:09:88:00:9a to interface 1.1 on VLAN internal.

```
show vlan
```

Displays all dynamic and static entries in the Layer 2 Forwarding table.

```
list vlan all-properties
```

Displays all properties for all static entries in the Layer 2 Forwarding table.

```
list vlan non-default-properties
```

Displays all non-default properties for all static entries in the Layer 2 Forwarding table.

```
delete vlan all
```

Deletes all entries in all VLAN Layer 2 Forwarding tables.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

all-records

Deletes, from the specified VLAN, all dynamic and static records.

description

User defined description.

dynamic

Displays or deletes all dynamic entries in the Layer 2 Forwarding table.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interface

Specifies an interface to which you want to map a MAC address. You must specify either an interface or a trunk when you create an entry in the Layer 2 Forwarding table.

MAC address

Specifies a 6-byte ethernet address in not case-sensitive hexadecimal colon notation, for example, 00:0b:09:88:00:9a. You must specify a MAC address when you create an entry in the Layer 2 Forwarding table.

partition

Displays the administrative partition in which the VLAN resides.

records

Specifies MAC addresses for the VLAN Layer 2 Forwarding table. Specifies MAC addresses that you want to add to, delete from, modify, or replace in the VLAN Layer 2 Forwarding table.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

static

Displays or deletes all static entries in the Layer 2 Forwarding table.

trunk

Specifies a trunk to which you want to map a MAC address. You must specify either an interface or a trunk when you create an entry in the Layer 2 Forwarding table.

SEE ALSO

delete, edit, glob, list, modify, net vlan, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013. All rights reserved.

net ike-evt-stat

NAME

ike-evt-stat - Displays and resets IKE event statistics

MODULE

net

SYNTAX

Display and reset the ike-evt-stat component within the net module using the syntax in the following section.

MODIFY

reset-stats ike-evt-stat

DISPLAY

show ike-evt-stat

DESCRIPTION

You can use the ike-evt-stat component to display and reset IKE event statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2014-07-10 net ike-evt-stat(1)

net ike-msg-stat

NAME

ike-msg-stat - Displays and resets IKE message statistics

MODULE

net

SYNTAX

Display and reset the ike-msg-stat component within the net module using the syntax in the following section.

MODIFY

reset-stats ike-msg-stat

DISPLAY

show ike-msg-stat

DESCRIPTION

You can use the ike-msg-stat component to display and reset IKE message statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2014-07-10 net ike-msg-stat(1)

net interface-cos

NAME

interface-cos - Displays and resets COS (Class of Service) related statistics for the interfaces.

MODULE

net

SYNTAX

Display cos related statistics within the net module using the following syntax.

MODIFY

reset-stats interface-cos

DISPLAY

show interface-cos

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

global

DESCRIPTION

You can use the interface-cos component to display COS related statistics, including pkts out and bits out for all 8 COS queue.

EXAMPLES

show interface-cos

Displays interface COS related statistics for the system.

For information about the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2015. All rights reserved.

BIG-IP 2015-07-22 net interface-cos(1)

net interface-ddm

NAME

interface-ddm - Displays DDM (Digital Diagnostics Monitoring) related statistics for the optical interfaces.

MODULE

net

SYNTAX

Display DDM related statistics within the net module using the following syntax.

DISPLAY

show interface-ddm

DESCRIPTION

You can use the interface-DDM component to display DDM related statistics, to view laser transmit and receive optical power levels.

EXAMPLES

show interface-ddm

Displays interface DDM related statistics for the system.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-08-07 net interface-ddm(1)

net interface

NAME

interface - Configures the parameters of interfaces.

MODULE

net

SYNTAX

Configure the interface component within the net module using the syntax in the following sections.

MODIFY

modify interface [name]

options:

description [string]

[disabled | enabled]

flow-control (none |rx | tx | tx-rx)

force-gigabit-fiber [enabled | disabled]

forward-error-correction [enabled | disabled]

media [auto | 10baseT half | 10baseT full | 100baseTX half |

100baseTX full | 1000baseT half | 1000baseT full |

1000baseSX full | 1000baseLX full | 1000baseCX full |

10GbaseT full | 10GbaseSR full | 10GbaseLR full |

10GbaseER full | 10SFP+Cu full | 40GbaseSR4 full |

40GbaseLR4 full | no-phy]

media-fixed [auto | 10baseT half | 10baseT full |

100baseTX half | 100baseTX full | 1000baseT half |

1000baseT full | no-phy]

media-sfp [auto | 10baseT half | 10baseT full | 100baseTX half |

100baseTX full | 1000baseT half | 1000baseT full |

1000baseSX full | 1000baseLX full | 1000baseCX full |

10GbaseT full | 10GbaseSR full | 10GbaseLR full |

10GbaseER full | 10SFP+Cu full | 40GbaseSR4 full |

40GbaseLR4 full | 100GbaseSR4 full | 100GbaseLR4 full |

none | no-phy]

no-mgmt

port-fwd-mode [I3 | passive | virtual-wire]

prefer-port [fixed | sfp]

sflow {

options:

poll-interval [integer]

poll-interval-global [no | yes]

}

stp [disabled | enabled]

stp-auto-edge-port [enabled | disabled]

stp-edge-port [false | true]

stp-link-type [auto | p2p | shared]

stp-reset

span-mode [false | true]

qinq-ethertype [string]

bundle [disabled | enabled | not-supported]

bundle-speed [100G | 40G | not-supported]

edit interface [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats interface

reset-stats interface [[[name] | [glob] | [regex]] ...]

DISPLAY

list interface

list interface [[[name] | [glob] | [regex]] ...]

show running-config interface

show running-config interface

[[[name] | [glob] | [regex]] ...]

options:

all-properties

mac-address

media-active

media-capabilities

media-max

mtu

non-default-properties

(pending | not-pending)

one-line

show interface

```
show interface [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DESCRIPTION

You can use the interface component to enable or disable an interface, and to display and set media options, duplex mode, and status for an interface. In addition, you can specify whether the interface participates in the spanning tree protocol (STP) configuration, and set per-interface STP parameters such as link type, edge port status, and automatic edge port detection.

EXAMPLES

```
modify interface 1.1 enabled
```

Enables the interface named 1.1.

```
modify interface 1.2 disabled
```

Disables the interface named 1.2.

```
modify interface 1.1 1.2 1.3 stp disable
```

Disables STP on the interfaces named 1.1, 1.2, and 1.3.

```
modify interface 1.1 1.2 1.3 stp-auto-edge-port enabled
```

Sets auto edge detection for STP on the interfaces named 1.1, 1.2, and 1.3.

```
modify interface 1.1 1.2 1.3 stp-edge-port true
```

Sets the edge port attribute for STP on the interfaces named 1.1, 1.2, and 1.3.

OPTIONS

description

User defined description.

[disabled | enabled]

Enables or disables the specified interface. The default value is enabled.

flow-control

Specifies how the system controls the sending of PAUSE frames for flow control. The default value is tx-rx.

none Disables flow control.

rx Specifies that the interface honors pause frames from its partner, but does not generate pause frames.

tx Specifies that the interface ignores pause frames from its partner, and generates pause frames when necessary.

tx-rx

Specifies that the interface honors pause frames from its partner, and also generates pause frames when necessary.

force-gigabit-fiber

Enables or disables forcing of gigabit fiber media. If this is enabled for a gigabit fiber interface, the media setting will be forced, and no auto-negotiation will be performed. If it is disabled, auto-negotiation will be performed with just a single gigabit fiber option advertised.

forward-error-correction

Enables or disables IEEE 802.3bm Clause 91 Reed-Solomon Forward Error Correction (RS-FEC) on 100G interfaces. This setting is not valid for LR4 media. This setting must be the same on both ends of the link.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

if-index

Displays the index assigned to this interface. It is a unique identifier assigned for all objects displayed in the SNMP IF-MIB.

mac-address

Displays the 6-byte ethernet address in not case-sensitive hexadecimal colon notation, for example, 00:0b:09:88:00:9a.

media

Specifies the settings for the interface. The possible values are: 10baseT-full, 10baseT-half, 10GbaseER full, 10GbaseLR-full, 10GbaseSR-full, 10GbaseT-full, 10SFP+Cu-full, 40GbaseSR4-full, 40GbaseLR4-full, 100baseTX-half, 100baseTX-full, 1000baseLX full, 1000baseCX-full, 1000baseT-full, 1000baseT-half, 1000baseSX-full, auto, and no-phy.

When you set the media option, the system automatically sets either the media-sfp or media-fixed option, based on whether the interface supports small form factor pluggable (SFP) interfaces, or for combo ports whether SFP is the preferred port.

media-active

Displays the current media setting for the interface.

media-fixed

Specifies the settings for a fixed (non-pluggable) interface. Use this option only with a combo port to specify the media type for the fixed interface, when it is not the preferred port.

media-max

Displays the maximum media value for the interface.

media-sfp

Specifies the settings for an SFP (pluggable) interface. Note that you use this option only with a combo port to specify the media type for the SFP interface, when it is not the preferred port.

module-description

Displays an optics module description about the interface. It is only available on a SFP/SFP+/XFP/QSFP+/QSFP28 unit.

transmitter-technology

Displays the transmitter technology of the pluggable unit on an interface. It is only available on a QSFP+/QSFP28 unit.

mtu Displays the Maximum Transmission Unit (MTU) of the interface, which is the maximum number of bytes in a frame without IP fragmentation.

name Specifies an interface name, for example 3.1, where 3 is the physical slot number holding the network interface hardware and 1 is the physical port number of that interface on that hardware. Another example is mgmt, the name given to the management interface.

no-mgmt

Ensures that no changes are made to the mgmt interfaces when all is specified. This is especially convenient when disabling all traffic interfaces using the disabled command.

[pending | not-pending]

Pending indicates that the slot with which the interface is associated does not contain a blade. Not-pending indicates that the slot with which the interface is associated is not a cluster member. The default value is pending.

port-fwd-mode

Specifies the port forwarding mode. The default value is I3.

The options are:

I3 Specifies to use the I3 mode. This is the traditional mode of operation of BIGIP, wherein BIGIP forwards traffic between 2 VLANs.

passive

Specifies to use the passive mode. Use this mode to enable the BIGIP to process out of band traffic. In this mode, the BIGIP assumes the received traffic is a copy of real inline traffic, and while it processes it like regular traffic, traffic received on such an interface is never forwarded out.

virtual-wire

Specifies to use the virtual-wire mode. Use this mode to enable the BIGIP to operate like a L2 switch - transparent to the neighboring devices. Packets received on a virtual-wire interface can only be forwarded to another virtual-wire-configured interface. Such traffic retains the L2 header and is forwarded to the same VLAN.

prefer-port

Indicates which side of a combo port the interface uses, if both sides of the port have the potential for external links. The default value is sfp. Do not use this option for non-combo ports.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

sflow

Specifies sFlow settings for the interface:

poll-interval

Specifies the maximum interval in seconds between two pollings. The default value is 0. To enable this setting, you must also set the poll-interval-global setting to no.

poll-interval-global

Specifies whether the global interface poll-interval setting, which is available under sys sflow global-settings module, overrides the object-level poll-interval setting. The default value is yes.

The available values are:

no Specifies to use the object-level poll-interval setting.

yes Specifies to use the global interface poll-interval setting.

serial

Displays the serial number of the pluggable unit on an interface. It is only available on a SFP/SFP+/XFP/QSFP+/QSFP28 unit.

`stp` Enables or disables STP. If you disable STP, no STP, RSTP, or MSTP packets are transmitted or received on the interface or trunk, and spanning tree has no control over forwarding or learning on the port or the trunk. The default value is enabled.

`stp-auto-edge-port`

Sets the STP automatic edge port detection for the interface. The default value is enabled. When STP automatic edge port detection is set to enabled on an interface, the system monitors the interface for incoming STP, RSTP, or MSTP packets. If no such packets are received for a sufficient period of time (about three seconds), the interface is automatically given edge port status. When automatic edge port detection is set to disabled on an interface, the system does not automatically give the interface the edge port status. Any STP setting set on a per-interface basis applies to all spanning tree instances.

`stp-edge-port`

Sets STP edge port detection. The default value is true.

`stp-link-type`

Specifies the STP link type for the interface. The default value is auto.

The spanning tree system includes important optimizations that can only be used on point-to-point links. That is, on links that connect just two bridges. If these optimizations are used on shared links, incorrect or unstable behavior may result. By default, the implementation assumes that full-duplex links are point-to-point and that half-duplex links are shared.

The options are:

`auto` Specifies that the system determines the spanning tree link type, which is based on the duplex setting.

`p2p` Specifies that the system uses the optimizations for point-to-point spanning tree links. Point-to-point links connect only two spanning tree bridges.

`shared`

Specifies that the system uses the optimizations for shared spanning tree links. Shared links connect two or more spanning tree bridges.

`stp-reset`

Resets STP, which forces a migration check.

`qinq-ethertype`

Specifies the ether-type value used for the packets handled on this port when it is a member in a QinQ vlan. The ether-type can be set to any string containing a valid hexadecimal 16 bits number, or any of the well known ether-types: 0x8100, 0x9100, 0x88a8. Default value is set to 0x8100.

`bundle`

Sets the bundle capability on the port.

`bundle-speed`

Sets the bundle-speed on the port when bundle capability is enabled.

`if-index`

Displays the index assigned to this interface. It is a unique identifier assigned for all objects displayed in the SNMP IF-MIB.

`vendor`

Displays the name of the vendor of the pluggable unit on an interface. It is only available on a SFP/SFP+/XFP/QSFP+/QSFP28 unit.

`vendor-partnum`

Displays the part number programmed by the vendor of the pluggable unit on an interface. It is only available on a SFP/SFP+/XFP/QSFP+/QSFP28 unit.

`vendor-oui`

Displays the vendor OUI of the pluggable unit on an interface. It is only available on a SFP/SFP+/XFP/QSFP+/QSFP28 unit.

`vendor-revision`

Displays the vendor revision of the pluggable unit on an interface. It is only available on a SFP/SFP+/XFP/QSFP+/QSFP28 unit.

SEE ALSO

`edit`, `glob`, `list`, `modify`, `regex`, `reset-stats`, `show`, `tms`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

net ipsec-stat

NAME

action - Displays and resets IPsec statistics.

MODULE

ipsec-stat

SYNTAX

Display statistics for the action component within the IPsec module using the syntax in the following section.

DISPLAY

show action

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the action component to display IPsec statistics. The statistics details are described below:

Cmd Id

Specifies the CPU id for IPsec.

Mode Specifies the IPsec mode (TRANSPORT/TUNNEL).

Proto

Specifies the Ipsec Protocol (ESP/AH)..

Packets In

Specifies the number of incoming IPsec packets.

Bytes In

Specifies the number of incoming IPsec bytes.

Packets Out

Specifies the number of outgoing IPsec packets.

Bytes Out

Specifies the number of outgoing IPsec bytes.

EXAMPLES

show action

Displays IPsec statistics.

reset-stats action

Resets the IPsec statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2017. All rights reserved.

BIG-IP 2017-03-02 net ipsec-stat(1)

net ipsec ike-daemon

NAME

ike-daemon - Configures the Internet Key Exchange (ISAKMP) daemon.

MODULE

net ipsec

SYNTAX

Configure the ike-daemon component within the net ipsec module using the syntax in the following sections.

MODIFY

modify ike-daemon

options:
description [string]
isakmp-natt-port [port number]
isakmp-port [port number]
log-level [error|warning|notify|info|debug|debug2]
natt-keep-alive [seconds]
log-publisher [string]

DISPLAY
list
list ike-daemon
show running-config ike-daemon
options:
all-properties
non-default-properties
one-line

DESCRIPTION
You can use the ike-daemon component to configure global settings for the IKE agent.

EXAMPLES
modify ike-daemon isakmp-port 500

Sets the isakmp port to 500.

OPTIONS
description
User defined description.

isakmp-natt-port
Specifies the port that the IKE daemon uses to accept ISAKMP messages when NAT-Traversal is detected. This is also the port number used to accept UDP-encapsulated ESP traffic for NAT-Traversal. Only 4500 is currently supported.

isakmp-port
Specifies the port that the IKE daemon uses to accept ISAKMP messages. Only 500 is currently supported.

log-level
Specifies the logging level of the IKE daemon. The log file is located at /var/log/racoon.log.

natt-keep-alive
Specifies the interval between sending NAT-Traversal keep-alive packets. The default value is 20 seconds. Set to 0 to disable keep-alive packets.

log-publisher
Specifies the logging publisher. A new log-publisher object can be created via TMSH command tms create sys log-config publisher.

SEE ALSO
list, net ipsec ike-peer, tms

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2017-05-26 net ipsec ike-daemon(1)

net ipsec ike-peer

NAME
ike-peer - Configures one or more IKE peers for IPsec.

MODULE
net ipsec

SYNTAX
Configure the ike-peer component within the net ipsec module using the syntax in the following sections.

CREATE/MODIFY
create ike-peer [name]
modify ike-peer [name]
options:
app-service [[string] | none]
ca-cert-file [certificate file]
crl-file [CRL file]
description [string]

```

dpd-delay [integer]
generate-policy [off | on | unique ]
lifetime [minutes]
mode [main | aggressive]
my-cert-file [certificate file]
my-cert-key-file [certificate key file]
my-cert-key-passphrase [none | [string] ]
my-id-type [address | asn1dn | fqdn | keyid-tag | user-fqdn]
my-id-value [string]
nat-traversal [on | off | force]
ocsp-cert-validator [ocsp-cert-validator-name-string]
ocsp-lifetime [minutes]
ocsp-jitter-percent [zero-to-fifty-percent]
passive [true | false]
peers-cert-file [certificate file]
peers-cert-type [certfile | none]
peers-id-type [address | asn1dn | fqdn | keyid-tag | user-fqdn]
peers-id-value [string]
phase1-auth-method [pre-shared-key | rsa-signature | dss | ecdsa-256 | ecdsa-384 | ecdsa-521 ]
phase1-encrypt-algorithm [3des | aes | blowfish | camellia | cast128 | des]
phase1-hash-algorithm [md5 | sha1 | sha256 | sha384 | sha512]
phase1-perfect-forward-secrecy [modp1024 | modp1536 | modp2048 | modp3072 | modp4096 | modp6144 | modp768 | modp8192 | ecp25
preshared-key [string]
preshared-key-encrypted [string]
prf [sha1 | sha256 | sha384 | sha512]
proxy-support [disabled | enabled]
remote-address [ip address]
replay-window-size [integer]
state [disabled | enabled]
traffic-selector [name]
verify-cert [true | false]
version [add | delete | none | replace-all-with] {
    [v1|v2]
}

```

DISPLAY

```

list ike-peer
list ike-peer [name]
show running-config ike-peer
show running-config ike-peer [name]
options:
    all-properties
    non-default-properties
    one-line

```

DELETE

```

delete ike-peer
delete ike-peer [name]

```

DESCRIPTION

You can use the ike-peer component to modify the IKE phase 1 parameters for each remote IKE peer. The setting in the default anonymous ike-peer will apply to any peer that does not match a more specific ike-peer directive.

EXAMPLES

```
create ike-peer SanJose { remote-address 1.2.3.4 preshared-key abc phase1-auth-method pre-shared-key }
```

Creates an ike-peer named SanJose that has the IP address of 1.2.3.4 using preshared key as the authentication method.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

ca-cert-file

Specifies the file name, which contains the certificates of the trusted root and intermediate certificate authorities.

crl-file

Specifies the file name of the Certificate Revocation List.

description

User-defined description.

dpd-delay

This option activates the Dead Peer Detection (DPD) and sets the time (in seconds) allowed between two proofs of liveness requests. The default value is 30. When the value is set to 0, it means to disable DPD monitoring, but still negotiate DPD support.

generate_policy

This directive is for the responder. To use it, set passive to true so the IKE peer is only a responder.

If the responder does not have any policy in the Security Policy Database (SPD) during phase 2 negotiation, and the directive is set to on, then the racoon daemon chooses the first proposal in the Security Association (SA) payload from the initiator, and generates policy entries from the proposal. It is useful to negotiate with clients whose IP address is allocated dynamically. If an inappropriate policy

is installed into the responder's SPD by the initiator, other communications might fail due to a policy mismatch between the initiator and the responder. The initiator ignores this directive. The default value is off.

lifetime

Specifies the lifetime of an IKE SA that will be proposed in the phase 1 negotiations.

mode Specifies the exchange mode for phase 1 when racoon is the initiator, or the acceptable exchange mode when racoon is the responder.

my-cert-file

Specifies the name of my certificate file. The certificate type must match the phase1-auth-method value. Note that there are no default certificates for DSS and ECDSA authentication methods.

my-cert-key-file

Specifies the name of my certificate key file. The certificate key type must match the phase1-auth-method value. Note that there are no default keys for DSS and ECDSA authentication methods.

my-cert-key-passphrase

Specifies the passphrase of the key used for my-cert-key-file. Note that only IKEv2 supports passphrase.

my-id-type

Specifies the identifier type sent to the remote host to use in the phase 1 negotiation.

my-id-value

Specifies the identifier value sent to the remote host to use in the phase 1 negotiation.

nat-traversal

Enables use of the NAT-Traversal IPsec extension (NAT-T). NAT-T allows one or both peers to reside behind a NAT gateway (that is, performing address- or port-translation). The presence of NAT gateways along the path is discovered during the phase 1 handshake, and if found, NAT-T is negotiated. When NAT-T is in charge, all ESP and AH packets of a given connection are encapsulated into UDP datagrams (port 4500, by default). The options are:

force

NAT-T is used regardless of whether NAT is detected between the peers.

off NAT-T is not proposed/accepted. This is the default.

on NAT-T is used when a NAT gateway is detected between the peers.

passive

Specify true if you do not want to be the initiator of the IKE negotiation with this ike-peer.

peers-cert-file

Specifies the peer's certificate for authentication. Deprecated in IKEv2 configuration.

peers-cert-type

Specifies that the only peers-cert-type supported is certfile. Deprecated in IKEv2 configuration.

peers-id-type

Specifies that address, fqdn, asn1dn, user-fqdn, or keyid-tag can be used as peers-id-type.

peers-id-value

Specifies the peer's identifier to be received. If it is not defined, then the IKE agent will not verify the peer's identifier in the ID payload transmitted from the peer. The usage of peers-id-type and peers-id-value is the same as my-id-type and my-id-value except that the individual component values of an asn1dn identifier may be specified as * to match any value (for example, "C=XX, O=MyOrg, OU=*, CN=Mine").

phase1-auth-method

Defines the authentication method used for the phase 1 negotiation. Possible values are: pre-shared-key if using pre-shared-key, and rsa-signature, dss, ecdsa-256, ecdsa-384 or ecdsa-521 if using X.509 certificate-based authentication. Note that dss, ecdsa certificates are supported in IKEv2 only."

phase1-encrypt-algorithm

Specifies the encryption algorithm used for the ISAKMP phase 1 negotiation. This directive must be defined. Possible value is one of following: des, 3des, blowfish, cast128, aes, or camellia for Oakley.

phase1-hash-algorithm

Defines the hash algorithm used for the ISAKMP phase 1 negotiation. This directive must be defined. The algorithm should be one of following: md5, sha1, sha256, sha384, or sha512 for Oakley.

phase1-perfect-forward-secrecy

Defines the Diffie-Hellman group for key exchange to provide perfect forward secrecy. This directive must be defined in one of Diffie-Hellman groups: modp768, modp1024, modp1536, modp2048, modp3072, modp4096, modp6144 and modp8192, or one of Elliptic-Curve Diffie-Hellman groups: ecp256, ecp384 and ecp521. Note that ECDH is supported in IKEv2 only.

pre-shared-key

Specifies the pre-shared key for ISAKMP SAs. This field is valid only when phase1-auth-method is pre-shared-key.

pre-shared-key-encrypted

Specifies the pre-shared key for ISAKMP SAs. This field is valid only when phase1-auth-method is pre-shared-key. Stores pre-shared-key in encrypted form.

prf Specifies the pseudo-random function to derive keying material for all cryptographic operations.

proxy-support

If this value is enabled, both values of ID payloads in the phase 2 exchange are used as the addresses of end-point of IPsec-SAs. This attribute must be enabled, which is the default value. This field is used only for IKEv1.

remote-address

Specifies the IP address of the IKE remote node. The format required for specifying a route domain ID in an IP address is A.B.C.D%ID. For example, A.B.C.D%2, where the IP address A.B.C.D pertains to route domain 2. The route domain id should be same as the route domain id specified in the source/destination address of the traffic selector associated with this remote node.

replay-window-size

Specifies the replay window size of the IPsec SAs negotiated with the IKE remote node. This window limits the number of out-of-order IPsec packets that can be received relative to the packet with the highest sequence number that has been authenticated so far. Packets with older sequence numbers that are outside this range are rejected. The default value is 64. The valid range is from 4 to 255.

state

Enables or disables this IKE remote node.

traffic-selector

Specifies the names of the traffic-selector objects associated with this ike-peer.

verify-cert If set to true, the identifier sent by the remote host (as specified in its `my_identifier` statement) is compared with the credentials in the certificate as follows: Type `asn1dn`: the entire certificate subject name is compared with the identifier, e.g. `"C=XX, O=YY, ..."`. Type `address`, `fqdn`, or `user_fqdn`: The certificate's `subjectAltName` is compared with the identifier. If the two do not match, the negotiation will fail. The default value is false, which is not to verify the identifier using the peer's certificate.

version

Specifies which version of IKE to be used. The default value is `v1`. The following versions are available:

`v1` Specifies version IKEv1 will be used.

`v2` Specifies version IKEv2 will be used.

SEE ALSO

`create`, `modify`, `delete`, `list`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2019-02-15 net ipsec ike-peer(1)

net ipsec ike-sa

NAME

`ike-sa` - Displays IKE security associations on the BIG-IP(r) system.

MODULE

`net ipsec`

SYNTAX

Use the `ike-sa` component within the `ipsec` module to manage IKE security associations using the following syntax.

DISPLAY

`show ike-sa`

option:

`all-properties`

`peer-ip [IP address]`

`peer-name [name]`

`route-domain [integer]`

`traffic-selector [name]`

DESCRIPTION

You can use the `ike-sa` component to display information about IKE security associations in the system.

EXAMPLES

`show ike-sa all-properties`

Display detail information about IKE security associations.

OPTIONS

`peer-ip`

Specifies the peer IP address of the security associations that you want to display.

peer-name

Specifies the peer name of the security associations that you want to display.

route-domain

Specifies route domain used for traffic that you want to display. The default value is the default route domain.

traffic-selector

Specifies the name of the traffic-selector associated with the security associations that you want to display.

SEE ALSO

show

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2015-01-21 net ipsec ike-sa(1)

net ipsec ipsec-policy

NAME

ipsec-policy - Configures the IPsec security policy.

MODULE

net ipsec

SYNTAX

Configure the ipsec-policy component within the net ipsec module using the syntax in the following sections.

CREATE/MODIFY

create ipsec-policy [name]

modify ipsec-policy [name]

options:

app-service [[string] | none]

description [string]

ike-phase2-auth-algorithm [aes-gcm128 | aes-gcm192 | aes-gcm256 | aes-gmac128 | aes-gmac192 | aes-gmac256 | sha1 | sha256 | sha384]

ike-phase2-encrypt-algorithm [3des | aes128 | aes192 | aes256 | aes-gcm128 | aes-gcm192 | aes-gcm256 | aes-gmac128 | aes-gmac192 | aes-gmac256]

ike-phase2-lifetime [integer]

ike-phase2-lifetime-kilobytes [integer]

ike-phase2-perfect-forward-secrecy [modp1024 | modp1536 | modp2048 | modp3072 | modp4096 | modp6144 | modp768 | modp8192]

ipcomp [deflate | none | null]

mode [transport | tunnel | interface]

protocol [esp]

tunnel-local-address [ip address]

tunnel-remote-address [ip address]

DISPLAY

list ipsec-policy

list ipsec-policy [[name] | [glob] | [regex]] ...]

show running-config ipsec-policy

show running-config ipsec-policy [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

partition

DESCRIPTION

An ipsec-policy indicates the ipsec rule and action to be applied to the packets matched by the traffic-selector associated with this ipsec-policy.

EXAMPLES

```
create ipsec ipsec-policy tunnel_policy_sjc_sea { description "ipsec policy for the sjc-sea ipsec tunnel" mode tunnel tunnel-local-address 1.1.1.1 tunnel-remote-address 2.2.2.2 }
```

Creates the tunnel mode ipsec-policy tunnel_policy_sjc_sea.

```
delete ipsec ipsec-policy tunnel_policy_sjc_sea
```

Deletes the ipsec-policy tunnel_policy_sjc_sea.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description
User defined description.

ike-phase2-auth-algorithm
Specifies a payload authentication algorithm for ESP. This attribute is only valid when IKE is used to negotiate Security Associations. The possible options are: aes-gcm128, aes-gcm192, aes-gcm256, aes-gmac128, aes-gmac192, aes-gmac256, sha256, sha384, sha512 and sha1. The default value is aes-gcm128.

Note: Because aes-gcm and aes-gmac are authenticated encryption algorithms, when ike-phase2-auth-algorithm is set to aes-gcm or aes-gmac, ike-phase2-encrypt-algorithm has to be set to the identical algorithm with the same key length. sha256, sha384, sha512 and sha1 can only be used with an encryption algorithm that is NOT an authenticated encryption algorithm.

ike-phase2-encrypt-algorithm
Specifies an encryption algorithm for ESP. This attribute is only valid when IKE is used to negotiate security associations. The default value is aes-gcm128.

Note: Because aes-gcm and aes-gmac are authenticated encryption algorithms, when ike-phase2-encrypt-algorithm is set to one of these algorithms, ike-phase2-auth-algorithm has to be set to the identical algorithm with the same key length.

ike-phase2-lifetime
Specifies the lifetime duration in minutes, for the dynamically-negotiated security associations (SA). This attribute is only valid when IKE is used to negotiate security associations.

ike-phase2-lifetime-kilobytes
Specifies the lifetime duration in kilobytes, for the dynamically-negotiated security associations (SA). This attribute is only valid when IKE is used to negotiate security associations. A value of '0' means the SA will not re-key based on the number of bytes encrypted/decrypted. The minimum recommended value is 1000 kilobytes. This value is not negotiated between peers."

ike-phase2-perfect-forward-secrecy
Defines the group of Diffie-Hellman exponentiations. This attribute is only valid when IKE is used to negotiate Security Associations. The value 'none' indicates that the PFS is disabled for phase2 SA negotiations.

mode Specifies a security protocol mode for use. The options are:

transport
IPsec transport mode is used.

tunnel
IPsec tunnel mode is used.

interface
IPsec interface mode is used.

protocol
Specifies the IPsec protocol: Encapsulating Security Payload (ESP) or Authentication Header (AH).

ipcomp
Specifies the compression algorithm for IPComp. The following codec are available:

none Disable IPComp

deflate
Packets will be encapsulated with IPComp header and Deflate compression algorithm will be applied to the data.

null Packets will be encapsulated with IPComp header but no compression algorithm will be applied to the data.

tunnel-local-address
Specifies the IP address of the local IPsec tunnel endpoint. This option is only valid when mode is tunnel. The format required for specifying a route domain ID in an IP address is A.B.C.D%ID. For example, A.B.C.D%2, where the IP address A.B.C.D pertains to route domain 2.

tunnel-remote-address
Specifies the IP address of the remote IPsec tunnel endpoint. This option is only valid when mode is tunnel. The format required for specifying a route domain ID in an IP address is A.B.C.D%ID. For example, A.B.C.D%2, where the IP address A.B.C.D pertains to route domain 2.

SEE ALSO

list, net ipsec traffic-selector, net ipsec manual-security-association, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

net ipsec ipsec-sa

NAME

ipsec-sa - Displays IPsec security associations on the BIG-IP(r) system.

MODULE

net ipsec

SYNTAX

Use the ipsec-sa component within the ipsec module to manage IPsec security associations using the following syntax.

DISPLAY

show ipsec-sa

option:

all-properties

src-addr [IP address]

dst-addr [IP address]

route-domain [integer]

spi [integer]

traffic-selector [name]

DESCRIPTION

You can use the ipsec-sa component to display information about IPsec security associations in the system.

EXAMPLES

show ipsec-sa all-properties

Display detail information about IPsec security associations.

OPTIONS

src-addr

Specifies the source IP address of the security associations that you want to display.

dst-addr

Specifies the destination IP address of the security associations that you want to display.

route-domain

Specifies route domain used for traffic that you want to display. The default value is the default route domain.

spi Specifies the SPI of the security associations that you want to display.

traffic-selector

Specifies the name of the traffic-selector object associated with the security associations that you want to display.

SEE ALSO

show, net ipsec traffic-selector, net ipsec ipsec-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2015-01-21 net ipsec ipsec-sa(1)

net ipsec manual-security-association

NAME

manual-security-association - Configures the IPsec manual-security-association.

MODULE

net ipsec

SYNTAX

Configure the manual-security-association component within the net ipsec module using the syntax in the following sections.

CREATE/MODIFY

```
create manual-security-association
modify manual-security-association
options:
  app-service [[string] | none]
  description [string]
  auth-algorithm [sha1]
  auth-key [key]
  destination-address [ip address]
  encrypt-algorithm [3des|aes128|aes192|aes256|null]
  encrypt-key [key]
  ipsec-policy [name]
  protocol [esp]
  source-address [ip address]
  spi [number]
```

DISPLAY

```
list manual-security-association
show running-config manual-security-association
options:
  app-service
  all-properties
  non-default-properties
  one-line
```

DELETE

```
delete manual-security-association [name]
```

DESCRIPTION

Manually configures Security Association Database(SAD) entries. Because each SA provides data protection only for unidirectional traffic, you must configure a manual-security-association for traffic in each direction to establish a bidirectional IPsec tunnel.

EXAMPLES

```
create ipsec manual-security-association msa_on_dut2_transport_in { auth-key test description "manual security
association on dut2 for dut1 - transport" destination-address 7.7.7.7 encrypt-key test ipsec-policy
transport_policy_on_dut2 source-address 2.2.2.2 spi 1025 }
```

Creates a manual-security-association object named msa_on_dut2_transport_in to use IPsec to protect traffic from 2.2.2.2 to 7.7.7.7 with the authentication key test and the encryption key test. The ipsec-policy object named transport_policy_on_dut2 is associated with this manually configured security association.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

auth-algorithm

Specifies an authentication algorithm.

auth-key

Specifies the key for the authentication algorithm.

auth-key-encrypted

Displays the encrypted auth-key.

description

User-defined description.

destination-address

Specifies the destination of the security association.

encrypt-algorithm

Specifies an encryption algorithm.

encrypt-key

Specifies the key for the encryption algorithm.

encrypt-key-encrypted

Display the encrypted encrypt-key.

ipsec-policy

Specifies the ipsec-policy associated with this manual-security-association.

protocol

Specifies the IPsec protocol: Encapsulating Security Payload (ESP) or Authentication Header (AH).

source-address

Specifies the source address of the security association.

spi Specifies the Security Parameters Index. If this is the Security Association(SA) for the outbound traffic, make sure it matches the SPI of the inbound SA configured on the remote site and vice versa. SPI values between 0 and 255 are reserved for the future use by IANA and cannot be used.

SEE ALSO

list, net ipsec ipsec-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2017-03-31 net ipsec manual-security-association(1)

net ipsec traffic-selector

NAME

traffic-selector - Configures a traffic selector for IPsec.

MODULE

net ipsec

SYNTAX

Configure the traffic-selector component within the net ipsec module using the syntax in the following sections.

CREATE/MODIFY

create traffic-selector [name]

modify traffic-selector [name]

options:

action [protect]

app-service [[string] | none]

description [string]

destination-address [ip address/netmask]

destination-port [port number]

direction [both | in | out]

ipsec-policy [name]

ip-protocol [protocol number]

order [integer]

source-address [ip address/netmask]

source-port [port number]

DISPLAY

list traffic-selector

list traffic-selector [name]

show traffic-selector

show traffic-selector [name]

DELETE

delete traffic-selector

delete traffic-selector [name]

DESCRIPTION

You can use the traffic-selector component to specify which incoming traffic you want the system to protect with IPsec.

EXAMPLES

```
create traffic-selector sjc2sea { source-address 10.10.10.0/24 destination address 20.20.20.0/24 direction
both ipsec-policy my_policy }
```

Creates a traffic-selector named sjc2sea, which has the IP address of 10.10.10.0/24 using ipsec-policy named my_policy.

OPTIONS

action

Specifies how the system handles traffic that matches the criteria in the traffic selector. Only protect is currently supported.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

destination-address

Specifies the destination IP address of the traffic to be matched. The format required for specifying a route domain ID in an IP address is A.B.C.D%ID. For example, A.B.C.D%2, where the IP address A.B.C.D pertains to route domain 2.

destination-port

Specifies the destination port number of the traffic to be matched.

direction

Specifies the direction of traffic to be protected with IPsec. If the direction is both, use source-address and destination-address with respect to the outbound direction. The default value is both.

ip-protocol

Specifies the IP protocol of the traffic to be matched.

ipsec-policy

Specifies the name of the IPsec policy to be enforced on the matched traffic.

order

Specifies the order in which traffic is matched, if traffic can be matched to multiple traffic selectors. Traffic is matched to the the traffic selector with the highest priority (lowest order number).

source-address

Specifies the source IP address of the traffic to be matched. The format required for specifying a route domain ID in an IP address is A.B.C.D%ID. For example, A.B.C.D%2, where the IP address A.B.C.D pertains to route domain 2.

source-port

Specifies the source port number of the traffic to be matched.

SEE ALSO

list, net ipsec ipsec-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2012. All rights reserved.

BIG-IP 2017-03-09 net ipsec traffic-selector(1)

net ipv6-subscriber-prefix-length

NAME

ipv6-subscriber-prefix-length - Configure global IPV6 Subscriber Prefix Length.

MODULE

net

SYNTAX

Configure the ipv6-subscriber-prefix-length component within the net module using the syntax shown in the following sections.

MODIFY

```
modify ipv6-subscriber-prefix-length
options:
value [1 - 128, only multiples of 8]
```

```
edit ipv6-subscriber-prefix-length
```

```
options:
all-properties
non-default-properties
```

DISPLAY

```
list ipv6-subscriber-prefix-length
show running-config ipv6-subscriber-prefix-length
options:
all-properties
non-default-properties
one-line
```

DESCRIPTION

Provides the ability to configure global IPV6 Subscriber Prefix Length.

value

Specifies the Prefix Length value between 1 and 128 (Only multiples of 8 are allowed).

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2015. All rights reserved.

net lldp-globals

NAME

lldp-globals - configure lldp global settings.

MODULE

net

SYNTAX

Configure the lldp-globals component within the net module using the syntax shown in the following sections.

MODIFY

modify lldp-globals

options:

[disabled | enabled]

max-neighbors-per-port[integer]

reinit-delay[integer]

tx-delay[integer]

tx-hold[integer]

tx-interval[integer]

edit lldp-globals

options:

all-properties

non-default-properties

DISPLAY

list lldp-globals

options:

all-properties

non-default-properties

one-line

DESCRIPTION

Provides the ability to configure global LLDP settings.

[disabled | enabled]

To disable or enable LLDP feature on BIGIP.

max-neighbors-per-port

Specifies maximum neighbors per port BIGIP LLDP feature supports.

reinit-delay

Specifies time delay, in seconds, from when LLDP ports is disabled to a new LLDP initialization.

tx-delay

Specifies the minimum time delay, in seconds, between successive LLDP frame transmissions.

tx-hold

Specifies the multiplier to use to calculate the TTL of transmitted LLDP frame.

tx-interval

Specifies the time interval, in seconds, between LLDP frame transmission to its neighbors.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2016. All rights reserved.

net lldp-neighbors

NAME

lldp-neighbors - show lldp neighbors information.

MODULE
net

SYNTAX
show lldp-neighbors information within the net module using the syntax shown in the following sections.

DISPLAY
show lldp-neighbors
options:
all
all-properties

DESCRIPTION
show lldp neighbor information.

SEE ALSO
modify, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2016. All rights reserved.

BIG-IP 2017-01-20 net lldp-neighbors(1)

net mroute

NAME
mroute - Displays the multicast routes in the BIG-IP(r) system.

MODULE
net

SYNTAX
Display multicast route (mroute) entries within the net module using the syntax given below.

DISPLAY
show net mroute
options:
source [IP address]
group [multicast group address]

DESCRIPTION
Use the mroute component to display all the multicast routes in the BIG-IP(r) system. Results can be filtered using source option, group option, or both.

EXAMPLES
show net mroute

Displays all the multicast routes in the system.

show net mroute source 10.10.10.1

Displays all the multicast routes in the system whose source IP address is 10.10.10.1.

show net mroute group 224.1.0.13

Displays all the multicast routes in the system whose group multicast address is 224.1.0.13.

show net mroute source 10.10.10.1 group 224.1.0.13

Displays all the multicast routes in the system whose source IP address is 10.10.10.1 and group multicast address is 224.1.0.13.

OPTIONS
source
Specifies the source IP address of the multicast routes that you want to display.

group
Specifies the multicast group address to display. Only a single group may be specified with the group option.

SEE ALSO
show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2015. All rights reserved.

BIG-IP 2016-03-14 net mroute(1)

net multicast-globals

NAME

multicast-globals - Global IP multicast configuration options

MODULE

net

SYNTAX

Configure the multicast-globals component within the net module using the syntax shown in the following sections.

MODIFY

modify multicast-globals

options:

route-lookup-timeout [integer value: 0 ~ 2147483647]
max-pending-routes [integer value: 0 ~ 2147483647]
max-pending-packets [integer value: 0 ~ 2147483647]
rate-limit [disabled | enabled]

edit multicast-globals

options:

all-properties
non-default-properties

DISPLAY

list multicast-globals

show running-config multicast-globals

options:

all-properties
non-default-properties
one-line

DESCRIPTION

Configure global options related to IP multicast traffic processing.

EXAMPLES

modify multicast-globals route-lookup-timeout 2

Specifies the maximum lifetime, in seconds, of an incomplete MFC entry. The default value is 2.

modify multicast-globals max-pending-routes 256

Specifies the maximum number of incomplete MFC entries any TMM instance can have at one time. The default value is 256.

modify multicast-globals max-pending-packets 16

Specifies the maximum number of packets queued on behalf of a single incomplete MFC entry. The default value is 16.

modify multicast-globals rate-limit enabled

When enabled, the packet rate limit configured in the DB variable switchboard.maxmcastrate is enforced. Otherwise multicast packets are not rate limited. The default value is enabled.

OPTIONS

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2015-2016. All rights reserved.

BIG-IP 2017-05-24 net multicast-globals(1)

net ndp

NAME

ndp - Configures IPv6-to-Ethernet neighbor discovery display and control.

MODULE

net

SYNTAX

Configure the ndp component within the net module using the syntax in the following sections.

CREATE/MODIFY

create ndp [name]

options:

description [string]

ip-address [ip address]

mac-address [MAC address]

modify ndp [name]

options:

description [string]

mac-address [MAC address]

edit ndp [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ndp

list ndp [[[name] | [glob] | [regex]] ...]

show running-config ndp

show running-config ndp

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show ndp

show ndp [[[name] | [glob] | [regex]] ...]

options:

[all | dynamic | field-fmt | static]

DELETE

delete ndp [[all] | [name]...]

DESCRIPTION

Configures the IPv6-to-Ethernet address translation tables used by the IPv6 neighbor discovery protocol.

EXAMPLES

```
create ndp myNdp ip-address fec0:f515::c001 mac-address 00:0B:DB:3F:F6:57
```

Maps the IPv6 address, fec0:f515::c001, to the MAC address, 00:0B:DB:3F:F6:57, and the name of this entry is myNdp. Alternatively, the address can be used as the name, like the following example.

```
create ndp fec0:f515::c001 mac-address 00:0B:DB:3F:F6:57
```

Maps the IPv6 address, fec0:f515::c001, to the MAC address, 00:0B:DB:3F:F6:57.

```
show ndp
```

Displays all static and dynamic IPv6 address-to-MAC address mappings.

OPTIONS

ip-address

The IP address that is to be mapped. This is optional, and if not present, the name needs to be a string that represents a valid IP address.

description

User defined description.

dynamic

Displays dynamic IPv6 address-to-MAC address mapping.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

ip-address

The IP address that is to be mapped. This is optional, and if not present, the name needs to be a string that represents a valid IP address.

mac-address

Specifies a 6-byte Ethernet address in hexadecimal colon notation that is not case-sensitive, for example, 00:0b:09:88:00:9a.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

static

Displays static IPv6 address-to-MAC address mapping.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2014-03-18 net ndp(1)

net packet-filter-trusted

NAME

packet-filter-trusted - Modifies or displays trusted allow lists for packet filters.

MODULE

net

SYNTAX

Configure the packet-filter-trusted component within the net module using the syntax in the following sections.

MODIFY

modify packet-filter-trusted

options:

description [string]

ip-addresses none

ip-addresses

[add | delete | replace-all-with] {

[ip address ...]

}

mac-addresses none

mac-addresses

[add | delete | replace-all-with] {

[MAC address ...]

}

vlan none

vlan

[add | delete | replace-all-with] {

[vlan name ...]

}

edit packet-filter-trusted

DISPLAY

list packet-filter-trusted

show running-config packet-filter-trusted

options:

all-properties

non-default-properties

one-line

DESCRIPTION

Use the packet-filter-trusted component to create a layer of security for the traffic management system using trusted allow lists.

Trusted allow lists are lists of IP addresses, MAC addresses, and VLANs that are exempt from packet filter rules.

Important: By default, packet filtering is disabled. You must enable packet filtering using the Configuration utility. For more information, see the TMOS(r) Management Guide for BIG-IP(r) Systems.

EXAMPLE

Creates a trusted allow list that allows anything listed to bypass the packet filter.

In the following example, you have an administrative laptop that you want to have unrestricted access to the traffic management system. This is a laptop, and therefore it might have a different IP address from time to time. One way to solve the problem is to add a trusted MAC address. This trusted allow list example shows the laptop MAC address as 00:02:3F:3E:2F:FE. Now the laptop can access the traffic management system regardless of what address it boots with or to which VLAN it is connected, as long as it is on the same physical segment as the traffic management system.

Also in this example, the traffic management system is configured for basic firewalling of the private/internal network. This example shows a way to filter incoming traffic and allow outgoing traffic to be unrestricted. To do this, you add trusted VLANs that represent all traffic that originated on the internal network. Another way to do this is to use trusted IP addresses instead, for example, 192.168.26.0/24.

```
modify packet-filter-trusted {
  vlans add { internal1 internal2 }
  mac-addresses add { 00:02:3F:3E:2F:FE }
}
```

OPTIONS

description
User defined description.

ip-addresses
Specifies a list of source IP addresses. Any traffic matching a source IP address in the list is automatically allowed. This simplifies configuration of the packet filter to allow trusted internal traffic to be passed from VLAN to VLAN without a filter rule, including out to the Internet. Processing of traffic by this option occurs before rule list evaluation, making it impossible to override this option and mask out (block) certain types of traffic with a packet filter rule. This option is empty by default.

mac-addresses
Specifies a list of MAC addresses. The system allows any traffic matching a MAC address in the source address list. This simplifies configuration of the packet filter to allow trusted internal traffic to be passed from VLAN to VLAN without a filter rule, including out to the Internet. Processing of traffic by this option occurs before rule list evaluation, making it impossible to override this option and mask out (block) certain types of traffic with a packet filter rule. This option is empty by default.

vlans
Specifies a list of ingress VLANs. Any traffic received on a VLAN that is on the ingress VLAN list is automatically allowed. This simplifies configuration of the packet filter to allow trusted internal traffic to be passed from VLAN to VLAN without a filter rule, including out to the Internet. Processing of traffic by this option occurs before rule list evaluation, making it impossible to override this option and mask out (block) certain types of traffic with a packet filter rule. This option is empty by default.

SEE ALSO

edit, list, ltm virtual, modify, net packet-filter, net vlan, net vlan-group, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 net packet-filter-trusted(1)

net packet-filter

NAME

packet-filter - Configures packet filter rules.

MODULE

net

SYNTAX

Configure the packet-filter component within the net module using the syntax in the following sections.

CREATE/MODIFY

```
create packet-filter [name]
```

```
modify packet-filter [name]
```

options:

```
action [accept | continue | discard | reject]
```

```
app-service [[string] | none]
```

```
description [string]
```

```
logging [enabled | disabled]
```

```
order [integer]
```

```
rate-class [name]
```

```
rule "[BPF expression]"
```

```
vlan [name]
```

```
edit packet-filter [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
all-properties  
non-default-properties
```

```
reset-stats packet-filter
```

```
reset-stats packet-filter  
[ [ [name] | [glob] | [regex] ] ... ]
```

```
DISPLAY
```

```
list packet-filter
```

```
list packet-filter
```

```
[ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config packet-filter
```

```
show running-config packet-filter
```

```
[ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
all-properties  
non-default-properties  
one-line
```

```
show packet-filter
```

```
show packet-filter [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)  
field-fmt
```

```
DELETE
```

```
delete packet-filter [ all | [name] ]
```

DESCRIPTION

You can use the packet-filter component to create a layer of security for the traffic management system using packet filter rules.

The BIG-IP(r) system packet filters are based on the Berkeley Software Design Packet Filter (BPF) architecture. Packet filter rules are composed of four mandatory attributes and three optional attributes. The mandatory attributes are name, order, action, and rule. The optional attributes are vlan, logging, and rate-class. The rule attribute you choose defines the BPF script to match for the rule.

Important: By default, packet filtering is disabled. You must enable packet filtering using the Configuration utility. For more information, see the TMOS(r) Management Guide for BIG-IP(r) Systems.

EXAMPLES

You can create a set of rules that specify what incoming traffic you want the system to accept and how to accept it. See the examples following.

Example 1: Block spoofed addresses

This example prevents private IP addresses from being accepted on a public VLAN. This is a way of ensuring that no one can spoof private IP addresses through the external VLAN of the system. In this example, the system logs when this happens:

```
create packet-filter spoof_blocker {  
order 5  
action discard  
vlan external  
logging enabled  
rule " (src net 172.19.255.0/24) "  
}
```

Example 2: Allow restricted management access

You can provide restricted SSH and HTTPS access to the traffic management system for management purposes, and keep a log of that access. Note: This not the same management access you can get through the management port/interface (mgmt); that interface is not affected by any packet filter configuration, and if that is the only way you want to allow access to your system, this configuration is not necessary.

In the first rule shown below, SSH is allowed access from a single fixed-address administrative workstation, and each access is logged. In the subsequent rule, browser-based Configuration utility access is allowed from two fixed-address administrative workstations; however, access is not logged.

```
create packet-filter management_ssh {  
order 10  
action accept  
logging enabled  
rule " (proto TCP) and (src host 172.19.254.10) and  
(dst port 22) "  
}
```

```
create packet-filter management_gui {  
order 15  
action accept  
rule " (proto TCP) and (src host 172.19.254.2 or  
src host 172.19.254.10) and (dst port 443) "  
}
```

Example 3: Allow access to all virtual servers

In this final example, you can verify that all of the virtual servers in your configuration are reachable

from the public network. This is critical if you have decided to use a default-deny policy. This example also shows how to rate shape all traffic to the virtual server IP address with a default rate class (that can be overridden by individual virtual servers or iRules(r) later).

Note: This example has a single virtual server IP, and it does not matter what port traffic is destined for. If you want to be more specific, you can specify each service port, as well (for example, HTTP, FTP, telnet).

```
create packet-filter virtuals {
order 20
action accept
vlan external
rate class root
rule " ( dst host 172.19.254.80 ) "
}
```

OPTIONS

You can use these options with the packet-filter component to create packet filter rules:

action
Specifies how the system handles a packet that matches the criteria in the packet filter rule. There is no default; you must specify a value when you create a packet filter rule.

The possible values are:

accept
Indicates that the system accepts the packet, and stops processing additional packet filter rules, if there are any.

continue
Indicates that the system acknowledges the packet for logging or statistical purposes, but makes no decision on how to handle the packet. The system continues to evaluate traffic matching a rule with the Continue action, starting with the next packet filter rule in the list.

discard
Indicates that the system drops the packet, and stops processing additional packet filter rules, if there are any.

reject
Indicates that the system drops the packet, and also sends a reject packet to the sender, indicating that the packet was refused.

app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description
User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

logging
Enables or disables packet filter logging. If you omit this value, no logging is performed.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

order
Specifies a sort order greater than 0 (zero). No two rules may have the same sort order. There is a single, global list of rules. Each rule in the list has a relative integer order. The system first evaluates the rule with the lowest order value, and then evaluates all other rules based on ascent of the order value assigned to each rule.

For example, if there are 5 rules, numbered 500, 100, 300, 200, 201; the rule evaluation order is 100, 200, 201, 300, 500.

The system compares each packet to be filtered against the list of rules in sequence, starting with the first. Evaluation of the rule list stops on the first match that has an action of accept, discard or reject. A match on a rule with an action of none does not stop further evaluation of the rule list; the system updates the statistics count and generates a log if the rule indicates it, but otherwise rule processing continues with the next rule in the list.

F5 Networks recommends that you sequence rules for effect and efficiency; generally this means:

-- Assign the lowest order to more specific rules, so that the system will evaluate those rules first.

-- The system evaluates one expression with multiple criteria more efficiently than multiple expressions each with a single criterion.

This option is required.

rate-class
Specifies the name of a rate class. The value is the name of any existing rate class. If omitted, no rate filter is applied.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rule Specifies the BPF expression to match. The rule is mandatory, however you can leave it empty. If empty, the packet filter rule matches all packets.

vlan Specifies the VLAN to which the packet filter rule applies. The value for this option is any VLAN name currently in existence. If you omit this value, the rule applies to all VLANs. If you do not provide a VLAN name when you create a packet-filter, the rule applies to all VLANs.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, net packet-filter-trusted, net vlan, net vlan-group, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-10-25 net packet-filter(1)

net packet-tester security

NAME

packet-tester - Shows if packet with given parameters passes through data path, which AFM policies and rules will be applied to the packet, if it will be dropped or not. This is shown for Dos in global and Virtual server context, IP Intelligence and ACL in global, route domain and listener context. You can only use the show command with this component.

MODULE

net

SYNTAX

```
show packet-tester security
  dest-addr [IP address]
  source-addr [IP address]
  dest-port [TCP/UDP port]
  source-port [TCP/UDP port]
  protocol [protocol]
  src-vlan [source vlan name]
  check-staged[enable/disable]
  trigger-log[enable/disable]
  ttl[1 to 255]
  syn[SYN TCP FLAG]
  ack[ACK TCP FLAG]
  rst[RST TCP FLAG]
  fin[FIN TCP FLAG]
  push[PUSH TCP FLAG]
  urg[URG TCP FLAG]
```

DESCRIPTION

With user provided VLAN, source/destination IP addresses, TCP/UDP ports and protocol, the command will craft a packet and insert into data path to match these parameters against user configured DOS, ACL rules and IP intelligence global, route domain, VIP/SelfIP context, and return which policies, rules applied and the final action taken on packet. Both IPv4 and IPv6 addresses and IP/UDP/TCP/SCTP protocols are supported. Detail option with provide which specific policy and rule will be applied to such a packet. This command can be used as a diagnostic tool to trouble-shoot BigIP AFM configuration problem. It provides a faster way to identify which AFM config will have impact to the specified packet stream.

EXAMPLES

```
[root@bigip208:Active:Standalone] rpm # tmsh -s -m show net packet-tester security dst-addr 41.41.41.41 dst-
port 80 src-addr 8.8.8.1 src-port 99 protocol udp src-vlan /Common/internal detail
```

```
*****
```

```
Packet Tester Data:
```

```
*****
```

```
Source IP/Port:8.8.8.1/99 Src Vlan /Common/internal
Destination IP/Port:41.41.41.41/80
Packet Protocol: udp
Packet Trace Option: Check Staged:Disable, Trigger Log:Disable
```

```
Stage:Device-DoS
Result: Allow, No Anomaly
Other Information
```

Dos Vector: UDP flood
Dos White list: No
Log Config:Disable

Stage:Device-IP Intelligence
Result: No Policy
Other Information
Policy Name: unset
Source Hit Type: No Match
Source Category: unset
Drop Source:No
Destination Hit Type: No Match
Destination Category: unset
Drop Destination:No
Log Config:Disable

Stage:Device-Access Control
Result: Allow
Other Information
Policy Name: /Common/policy1
Policy Type: Enforced
Rule Name: packet_test_udp_rule
Source FQDN: No-lookup
Destination FQDN: No-lookup
Source Geo: No-lookup
Dest Geo: No-lookup
iRule:unset
Log Config:Disable

Stage:Route Domain-IP Intelligence (/Common/0)
Result: No Policy
Other Information
Policy Name: unset
Source Hit Type: No Match
Source Category: unset
Drop Source:No
Destination Hit Type: No Match
Destination Category: unset
Drop Destination:No
Log Config:Disable

Stage:Route Domain-Access Control (/Common/0)
Result: Allow
Other Information
Policy Name: /Common/policy1
Policy Type: Enforced
Rule Name: packet_test_udp_rule
Source FQDN: No-lookup
Destination FQDN: No-lookup
Source Geo: No-lookup
Destination Geo: No-lookup
iRule:unset
Log Config:Disable

Stage:Listener-DoS (/Common/packet_test_catchall)
Result: No Policy
Other Information
Dos Profile Name: unset
Dos Vector: unset
Dos White list: No
Log Config:Disable

Stage:Listener-IP Intelligence (/Common/packet_test_catchall)
Result: No Policy
Other Information
Policy Name: unset
Source Hit Type: No Match
Source Category: unset
Drop Source:No
Destination Hit Type: No Match
Destination Category: unset
Drop Destination:No
Log Config:Disable

Stage:Listener-Access Control (/Common/packet_test_catchall)
Result: Allow
Other Information
Policy Name: /Common/policy1
Policy Type: Enforced
Rule Name: packet_test_udp_rule
Source FQDN: No-lookup
Destination FQDN: No-lookup
Source Geo: No-lookup
Destination Geo: No-lookup
iRule:unset
Log Config:Disable

Final Result
Source IP/Port:8.8.1/99 Src Vlan /Common/internal
Destination IP/Port:41.41.41.41/80
Packet Protocol: udp
Packet Trace Option: Check Staged:Disable, Trigger Log:Disable
Final Action : Allow
Total records returned: 1

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-09-13 net packet-tester security(1)

net port-list

NAME

port-list - A port list is a list of IP ports.

MODULE

net

SYNTAX

CREATE/MODIFY

```
create port-list [name]
```

```
modify port-list [[name] | all]
```

options:

```
ports [add | delete | modify | replace-all-with] {
```

```
  [ [port] ]
```

```
}
```

```
app-service [name]
```

```
description [string]
```

```
edit port-list [[name] | all]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list port-list [[name] | all | [property]]
```

```
show running-config port-list [[name] | all | [property]]
```

DELETE

```
delete port-list [[name] | all]
```

DESCRIPTION

You can use the port-list component to define reusable lists of ports.

EXAMPLES

```
create port-list plist1 ports add { 80 443 }
```

Creates a new port list, "plist1," with two ports.

```
modify port-list plist1 description "my port list"
```

Modifies the above port list with a description.

```
modify port-list plist1 ports add { 999 }
```

Modifies the same port list by adding port 999 to it.

```
list port-list plist1
```

```
net port-list plist1 {
```

```
  ports {
```

```
    description "my port list"
```

```
    999 { }
```

```
    http { }
```

```
    https { }
```

```
  }
```

```
}
```

Shows the modified port list.

OPTIONS

ports

Specifies a list of IP ports.

The next keyword specifies the action to take with the ports (add, delete, modify, or replace the current set of ports).

add Creates a new port list.

delete
Deletes the port(s) that you specify next, in curly braces ({}).

modify
Makes it possible to replace the optional description(s) for the port(s). You can specify a description in a nested set of curly braces after each port.

replace-all-with
Replaces the current set of IP ports with the port(s) that you specify next, in curly braces ({}).

app-service
Associates this port list with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers.

description
Is your description for this port list.

SEE ALSO

edit, list, modify, ltm virtual, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2017. All rights reserved.

BIG-IP 2017-12-04 net port-list(1)

net port-mirror

NAME
port-mirror - Configures interface (port) mirroring.

MODULE
net

SYNTAX
Configure the port-mirror component within the net module using the syntax in the following sections.

CREATE/MODIFY
create port-mirror [interface_name]
modify port-mirror [interface_name]
options:
 app-service [[string] | none]
 interfaces
 [add | delete | replace-all-with] {
 [interface_name ...]
 }
 interfaces [default | none]

edit port-mirror [[interface_name] | [glob] | [regex]] ...]
options:
 all-properties

DISPLAY
list port-mirror
list port-mirror
 [[interface_name] | [glob] | [regex]] ...]
show running-config port-mirror
show running-config port-mirror
 [[interface_name] | [glob] | [regex]] ...]
options:
 one-line

DELETE
delete port-mirror [interface_name]

DESCRIPTION

You can use the port-mirror component to mirror traffic from interfaces on a blade to other interfaces on the same blade or another blade.

EXAMPLES

```
create port-mirror 1/1.1 interfaces add 1/1.2 1/1.3 1/1.4
```

Creates a port mirror from interface 1.1 on blade 1 to interfaces 1.2, 1.3, 1.4 on the same blade. The system mirrors traffic from interfaces 1.2, 1.3, and 1.4 on blade 1 to the interface 1.1 on the same blade.

```
modify port-mirror 1/1.1 interfaces delete 1/1.3 1/1.4
```

Deletes interfaces 1.3 and 1.4 on blade 1 from the existing port mirror 1/1.1 on the same blade.

OPTION

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

interface_name

Specifies the name of the interface, for example, 1/1.1.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, "net interface", regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 net port-mirror(1)

net rate-shaping class

NAME

class - Configures a rate class.

MODULE

net rate-shaping

SYNTAX

Configure the class component within the net rate-shaping module using the syntax in the following sections.

CREATE/MODIFY

```
create class [name]
```

```
modify class [name]
```

options:

app-service [[string] | none]

ceiling [integer]

ceiling-percentage [integer]

description [string]

direction [any | to-client | to-server]

drop-policy [[custom drop policy name] | fred | red | tail]

max-burst [integer]

parent [class name]

queue [[custom queue name | pfifo | sfq]

rate [integer]

rate-percentage [integer]

shaping-policy [[custom shaping policy name] | none]

```
edit class [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties

non-default-properties

DISPLAY

```
list class
```

```
list class [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config class
```

```
show running-config class [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties
current-module
non-default-properties
one-line

show class
show class [[[name] | [glob] | [regex]] ...]
options:
current-module
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DELETE
delete class [all | [name]]

DESCRIPTION

You can use the class component to create a rate class. A rate class lets you specify shaping properties for a specific type of traffic, such as Layer 3 traffic that specifies a certain source, destination, or service. Specifically, a rate class defines the number of bits per second that the system accepts per flow and the number of packets in a queue.

You configure rate shaping by creating a class and then assigning the class to a packet filter, a virtual server, or from within an iRule. When you configure a class, you can associate another class with the class you are configuring using the parent option.

You can also associate drop policies, shaping policies, and queues with a class using the drop-policy, shaping-policy, and queue options of the class component. You can associate pre-configured drop policies and queues with the class, or you can create custom drop policies, queues, and shaping policies, and then associate them with the class.

Note that if you specify a value for the shaping-policy option of the class, the system automatically changes the ceiling-percentage, drop-policy, max-burst, queue, and rate-percentage options of the class to match the values in the specified shaping policy.

EXAMPLES

```
create class my_class rate 10
```

Creates a class named my_class with a rate of 10.

```
list class all-properties
```

Displays all of the properties of all of the classes.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

ceiling

Specifies (in bps) how far beyond the value specified for the rate option that traffic can flow. This number sets an absolute limit. No traffic can exceed this rate. The rate class might limit traffic throughput to the value of the rate option when there is high contention among siblings of a parent-child class hierarchy. The default value is the value of the rate option. The minimum value is 296 bps.

ceiling-percentage

Specifies the ceiling of the rate class as a percentage of the ceiling of the associated parent class. This option applies only to rate classes with an associated parent rate class. The default value is 0 (zero), which indicates that the class uses the value of the ceiling option.

description

User defined description.

direction

Specifies the direction of traffic to which the class is applied. The default value is any.

drop-policy

Specifies the name of a drop policy. You can use one of the pre-configured drop policies, or you can create a customized drop policy using the drop-policy component.

The default value is tail, which is the simplest drop policy. The pre-configured drop policies are:

fred Specifies that the system uses Flow-based Random Early Detection to decide whether to drop packets based on the aggressiveness of each flow.

red Specifies that the system uses Random Early Detection to determine whether to drop packets to maintain the average queue length within the specified range.

tail Specifies that the system drops all incoming packets when the queue is full.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

max-burst

Specifies the maximum number of bytes that traffic can burst beyond the value of the rate option. The traffic may not burst higher than the value of the ceiling option. The default value is 0 (zero).

name Specifies a unique name for the component. This option is required for the commands create, delete, and

modify.

parent

Associates another class with this class. The class you are configuring (which when you configure a parent class for it becomes a child class) can borrow bandwidth from the parent class. The parent class can use any of the unused bandwidth of the child class.

queue

Specifies the queuing method. The default value is sfq. The pre-configured options are:

pfifo

The Priority FIFO queuing method queues all traffic under a set of five sub-queues based on the Type of Service (TOS) field of the traffic. Four of the sub-queues correspond to the four possible TOS values (Minimum delay, Maximum throughput, Maximum reliability, and Minimum cost). The fifth sub-queue represents traffic with no TOS value. The Priority FIFO method processes these five sub-queues in a way that preserves the meaning of the TOS value as much as possible. For example, a packet with the TOS value of Minimum cost might yield dequeuing to a packet with the TOS value of Minimum delay.

sfq Stochastic Fair Queuing is a queuing method that further queues packets under a set of many FIFO sub-queues. Selecting a specific sub-queue is based on a hash of the flow address information. SFQ dequeues packets from the set of sub-queues in a Round Robin fashion. The overall effect is that fairness of dequeuing is achieved, because packets from one flow cannot occupy the queues at the exclusion of those of another flow.

Note that if you assign a shaping policy to the class, then the queuing discipline of the class becomes that specified in the shaping policy. If you do not assign a shaping policy to the class, the default queue is sfq.

rate Specifies the guaranteed throughput rate of the traffic handled by this rate class, in bits per second (bps).

rate-percentage

Specifies the rate of the rate class as a percentage of the ceiling of the associated parent class. This option applies only to rate classes with an associated parent rate class. The default value is 0 (zero), which specifies that the system uses the value of the rate option.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

shaping-policy

Specifies the name of a shaping policy. The default value is none.

Note that the system automatically changes the ceiling-percentage, drop-policy, max-burst, queue, and rate-percentage options of this class to match the values in the specified shaping policy.

SEE ALSO

create, delete, edit, glob, list, modify, net rate-shaping drop-policy, net rate-shaping queue, net rate-shaping shaping-policy, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 net rate-shaping class(1)

net rate-shaping color-policer

NAME

color-policer - Configures a policer, which can be applied to other configuration entities to meter or rate-shape traffic using color based classification of packets.

MODULE

net rate-shaping

SYNTAX

Configure the color-policer component within the net rate-shaping module using the syntax in the following sections.

CREATE

create color-policer [name]

modify color-policer [name]

options:

action [action]

app-service [[string] | none]

committed-burst-size [integer]
description [string]
committed-information-rate [integer]
excess-burst-size [integer]

DISPLAY

list color-policer

options:

all-properties

DELETE

delete color-policer [name]

DESCRIPTION

You can use a color-policer to create a metering/policing configuration to be applied to other configuration entities. For instance, to limit or track out of profile network traffic to a vCMP guest, a color-policer can be added to a vcmp traffic-profile and be applied to any number of vcmp guest objects.

Based on the configuration of color-policer attributes a packet may be counted as either green, yellow, or red. These categories can be tracked for accounting purposes, as well as allowing for out of profile (Red) traffic to be dropped as a rate-shaping or DOS protection technique.

EXAMPLES

list net color-policer

Lists the current configuration of all color-policers.

```
create net color-policer fiftyMbpsLimiter action drop-red committed-information-rate 50mbps committed-burst-size 10mb excess-burst-size 10mb
```

Configures a policer with a committed information rate (CIR) of 50 mbps committed burst size (CBS) of 10 mb, and excess burst size (EBS) of 10 mb. See individual descriptions of options below for more detailed semantics.

OPTIONS

action

One of drop-red or default. Drop-red means that a packet which exceeds the excess burst rate of the policer, and is marked red should be dropped immediately before even attempting to enqueue the packet to/from the guest. Default means that the default behavior for the resulting color packets should be taken, which may vary depending on additional system configuration, and load. For instance, a Red colored packet may be considered lowest priority with respect to QOS queues associated with the guest. Thus, allowing as much traffic through as is available, but increasing the drop probability for packets out of profile in the case that the system is under high enough data-plane traffic load.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

committed-burst-size

When the data rate exceeds the committed-information-rate then the committed-burst-size (by default specified in bytes) is the burst size below which a packet is marked green, and above which it may be marked yellow, or red.

committed-information-rate

The committed rate of data transfer that is to be given to a metered entity that the policer is associated with. Traffic is marked green as long as it stays below this rate or if bursts above this rate are smaller than the committed-burst-size.

description

User defined description.

excess-burst-size

An additional data burst size to be used on top of the committed-burst-size. A packet is marked yellow if it exceeds the committed-burst-size but not the additional excess-burst-size, and red if it exceeds the excess-burst-size.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2016. All rights reserved.

BIG-IP 2016-03-14 net rate-shaping color-policer(1)

net rate-shaping drop-policy

NAME

drop-policy - Configures a custom drop policy for use in rate shaping.

MODULE

net rate-shaping

SYNTAX

Configure the drop-policy component within the net rate-shaping module using the syntax in the following sections.

CREATE/MODIFY

create drop-policy [name]

modify drop-policy [name]

options:

app-service [[string] | none]

average-packet-size [integer]

description [string]

fred-max-active [integer]

fred-max-drop [integer]

fred-min-drop [integer]

inverse-weight [integer]

max-probability [integer]

max-threshold [integer]

min-threshold [integer]

red-hard-limit [integer]

type [fred | red | tail]

edit drop-policy [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list drop-policy

list drop-policy [[[name] | [glob] | [regex]] ...]

show running-config drop-policy

show running-config drop-policy [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete drop-policy [all | [name]]

DESCRIPTION

A drop policy tells the system when and how to drop packets when the traffic handling queue is full, if required. The system comes with three pre-configured drop policies: fred, red, and tail.

You can use the drop-policy component to create a custom drop policy, and then associate it with a class using the drop-policy option of the class component. For more information, see `net rate-shaping class`.

You can also associate a custom drop policy with a shaping policy using the drop-policy option of the shaping-policy component. For more information, see `net rate-shaping shaping-policy`.

EXAMPLES

```
create drop-policy my_dp
```

Creates a custom drop policy named my_dp.

```
list drop-policy all-properties
```

Displays all of the properties of all of the drop policies.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

average-packet-size

Specifies the average MTU (maximum transmission unit) size in the range of 0 to 10000 bytes. The default value is 0 (zero).

description

User defined description.

fred-max-active

Specifies the maximum number of flows that can be active for each queue. The range is 0 to 10000. The default value is 0 (zero), which disables active flow limitation.

fred-max-drop

Specifies a hard drop limit in the range of 0 to 400. The default value is 0 (zero). Setting this to a small value does not change the hard drop limit, but a higher number increases the limit.

fred-min-drop

Specifies a hard no drop limit in the range of 0 to 100. The default value is 0 (zero). Setting this to a large value prevents packets from being dropped.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

inverse-weight

Specifies the weight used to calculate the average queue length. Valid values are 0, 64, 128, 256, 512, and 1024. The default value is 0 (zero).

max-probability

Specifies the maximum percentage probability in the range of 0 to 100 according to which packets are dropped when the average queue length is between the minimum and maximum thresholds. The default value is 0 (zero).

max-threshold

Specifies the queue length above which the system drops packets. The default value is 0 (zero).

min-threshold

Specifies the queue length below which the system does not drop packets. The default value is 0 (zero).

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

red-hard-limit

Specifies the maximum queue size in bytes. Additional packets are dropped. The default value is 0 (zero).

This option applies only when the value of the type option is red.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

type Specifies the type of drop policy. The default value is tail.

The options are:

fred Specifies that the system uses Flow-based Random Early Detection to decide whether to drop packets based on the aggressiveness of each flow.

red Specifies that the system uses Random Early Detection to determine whether to drop packets to maintain the average queue length within the specified range.

tail Specifies that the system drops all incoming packets when the queue is full. This is the simplest drop policy.

Note that although you could create a drop policy based on tail, that is already the default value of the drop-policy option in both the shaping-policy and class components.

SEE ALSO

create, delete, edit, glob, list, modify, net rate-shaping class, net rate-shaping queue, net rate-shaping shaping-policy, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 net rate-shaping drop-policy(1)

net rate-shaping queue

NAME

queue - Configures a custom queuing method.

MODULE

net rate-shaping

SYNTAX

Configure the queue component within the net rate-shaping module using the syntax in the following sections.

CREATE/MODIFY

create queue [pfifo | sfq]

modify queue [all | pfifo | sfq]

options:
app-service [[string] | none]
description [string]
pfifo-max-size [integer]
pfifo-min-size [integer]
sfq-bucket-count [integer]
sfq-bucket-size [integer]
sfq-perturbation [integer]
type [pfifo | sfq]

edit queue [[[all | pfifo | sfq] | [glob] | [regex]] ...]

options:
all-properties
non-default-properties

DISPLAY

list queue
list queue [[[all | pfifo | sfq] | [glob] | [regex]] ...]

show running-config queue
show running-config queue
[[[all | pfifo | sfq] | [glob] | [regex]] ...]

options:
all-properties
non-default-properties
one-line

DELETE

delete queue [all | [name]]

DESCRIPTION

You can use the queue component to configure a custom queuing method.

EXAMPLES

create queue my_q type pfifo

Creates a pfifo type queue name my_q.

list queue all-properties

Displays all of the properties of all of the queue.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

pfifo-max-size

Specifies the size of the largest queue for the pfifo type only. The default value is 0 (zero). Valid units are bytes(default), eb, gb, k, kb, mb, pb, and tb.

pfifo-min-size

Specifies the size of the smallest queue for the pfifo type only. The default value is 0 (zero). Valid units are bytes(default), eb, gb, k, kb, mb, pb, and tb.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

sfq-bucket-count

Specifies the number of buckets into which the queue is divided when you are configuring the sfq type. Valid values are 0, 16, 32, 64, 128, 256, 512, and 1024. The default value is 0 (zero).

sfq-bucket-size

Specifies the bucket size for the sfq type. The default value is 0 (zero). Valid units are bytes(default), eb, gb, k, kb, mb, pb, and tb.

sfq-perturbation

Specifies the interval in seconds at which the system reconfigures the SFQ hash function. This option applies only to the sfq type. The default value is 0 (zero).

type Specifies the queue discipline this custom queue uses. The options are:

pfifo

The Priority FIFO queuing method queues all traffic under a set of five sub-queues based on the Type of Service (TOS) field of the traffic. Four of the sub-queues correspond to the four possible TOS values (Minimum delay, Maximum throughput, Maximum reliability, and Minimum cost). The fifth sub-

queue represents traffic with no TOS value. The Priority FIFO method processes these five sub-queues in a way that preserves the meaning of the TOS value as much as possible. For example, a packet with the TOS value of Minimum cost might yield dequeuing to a packet with the TOS value of Minimum delay.

sfq Stochastic Fair Queuing is a queuing method that further queues packets under a set of many FIFO sub-queues. Selecting a specific sub-queue is based on a hash of the flow address information. SFQ dequeues packets from the set of sub-queues in a Round Robin fashion. The overall effect is that fairness of dequeuing is achieved, because packets from one flow cannot occupy the queues at the exclusion of those of another flow.

SEE ALSO

create, delete, edit, glob, list, modify, net rate-shaping class, net rate-shaping drop-policy, net rate-shaping shaping-policy, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 net rate-shaping queue(1)

net rate-shaping shaping-policy

NAME

shaping-policy - Configures a custom rate shaping policy for traffic flow.

MODULE

net rate-shaping

SYNTAX

Configure the shaping-policy component within the net rate-shaping module using the syntax in the following sections.

CREATE/MODIFY

```
create shaping-policy [name]
modify shaping-policy [name]
options:
  app-service [[string] | none]
  ceiling-percentage [integer]
  description [string]
  drop-policy [ [name] | none]
  max-burst [integer]
  queue [ [name] | none]
  rate-percentage [integer]
```

```
edit shaping-policy [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list shaping-policy
list shaping-policy [ [ [name] | [glob] | [regex] ] ... ]
show running-config shaping-policy
show running-config shaping-policy [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

DELETE

```
delete shaping-policy [all | [name] ]
```

DESCRIPTION

You can use the shaping-policy component to create a custom rate shaping policy to handle traffic flow, and then associate the shaping policy with a class.

Note that if you specify a value for the shaping-policy option of a class, the system automatically changes the ceiling-percentage, drop-policy, max-burst, queue, and rate-percentage options of that class to match the values in the shaping policy.

EXAMPLES

```
create shaping-policy my_sp
```

Creates a shaping policy named my_sp.

```
list shaping policies all-properties
```

Displays all of the properties of all of the shaping policies.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

ceiling-percentage

Specifies the percentage of the value of the ceiling option specified for the parent associated with the class component to which this shaping policy is associated. The default value is 0 (zero).

description

User defined description.

drop-policy

Specifies the name of a drop policy for this traffic flow. The default value is none.

You can use one of the pre-configured drop policies, or you can create a customized drop-policy using the drop-policy component.

The pre-configured drop policies are:

fred Specifies that the system uses Flow-based Random Early Detection to decide whether to drop packets based on the aggressiveness of each flow.

red Specifies that the system uses Random Early Detection to determine whether to drop packets to maintain the average queue length within the specified range.

tail Specifies that the system drops all incoming packets when the queue is full. This is the simplest drop policy.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

max-burst

Specifies the maximum number of bytes that traffic is allowed to burst beyond the value of the rate option of the class component to which this shaping policy is associated. The default value is 0 (zero).

Valid units are byte, bytes(default), eb, gb, k, kb, mb, pb, and tb.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

queue

Specifies the queuing method for this traffic flow. The default value is none. You can create a customized queuing method using the queue component. For more information, see net rate-shaping queue.

The preconfigured queues are:

pfifo

The Priority FIFO queuing method queues all traffic under a set of five sub-queues based on the Type of Service (TOS) field of the traffic. Four of the sub-queues correspond to the four possible TOS values (Minimum delay, Maximum throughput, Maximum reliability, and Minimum cost). The fifth sub-queue represents traffic with no TOS value. The Priority FIFO method processes these five sub-queues in a way that preserves the meaning of the TOS value as much as possible. For example, a packet with the TOS value of Minimum cost might yield dequeuing to a packet with the TOS value of Minimum delay.

sfq Stochastic Fair Queuing is a queuing method that further queues packets under a set of many FIFO sub-queues. Selecting a specific sub-queue is based on a hash of the flow address information. SFQ dequeues packets from the set of sub-queues in a Round Robin fashion. The overall effect is that fairness of dequeuing is achieved, because packets from one flow cannot occupy the queues at the exclusion of those of another flow.

rate-percentage

Specifies the percentage of the value of the rate option of the parent, which is associated with the class component to which this shaping policy is associated, that is available for this traffic flow. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, net rate-shaping drop-policy, net rate-shaping queue, net rate-shaping shaping-policy, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

net route-domain

NAME

route-domain - Configures route-domains for traffic management.

MODULE

net

SYNTAX

Configure the route-domain component within the net module using the syntax in the following sections.

CREATE/MODIFY

create route-domain [[name] | none]

options:

id [integer]

modify route-domain [name]

options:

app-service [[string] | none]

bwc-policy [string]

connection-limit [integer]

description [string]

flow-eviction-policy [[eviction policy name] | none]

fw-enforced-policy [[policy_name] | none]

fw-staged-policy [[policy_name] | none]

id [integer]

parent [[name] | none]

security-nat-policy [[policy_name] | none]

service-policy [[policy_name] | none]

strict [disabled | enabled]

routing-protocol

[add | delete | replace-all-with] {

[protocol name] ...

}

vlangs

[add | delete | replace-all-with] {

[vlan name] ...

}

edit route-domain [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats route-domain [name]

fw-enforced-policy-rules { [rule name] }

fw-staged-policy-rules { [rule name] }

security-nat-rules { [rule name] }

options:

fw-context-stat

ip-intelligence-categories

port-misuse

DISPLAY

list route-domain

list route-domain [[name] | [glob] | [regex]] ...]

show running-config route-domain

show running-config route-domain

[[name] | [glob] | [regex]] ...]

options:

all-properties

one-line

non-default-properties

show route-domain [[[name] | [glob] | [regex]] ...]

options:

fw-context-stat

ip-intelligence-categories

port-misuse

DELETE

delete route-domain [name]

DESCRIPTION

Using route domains, you can assign the same IP address to more than one device on a network, as long as each instance of the IP address resides in a separate routing domain.

EXAMPLES

create route-domain myRouteDomain id 1 vlans add { my_vlan }

Creates a route domain named myRouteDomain with an ID of 1 that includes my_vlan.

list route-domain all-properties

Displays all properties of all route domains.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

bwc-policy

Configures the bandwidth control policy for the route-domain. If set, it will enforce a throughput policy for incoming network traffic.

connection-limit

Configures the connection limit for the route domain. If set to a value other than zero, this specifies the total number of open connections allowed on this route domain. The default value is 0, unlimited.

description

User defined description.

id Specifies a unique numeric identifier for the route-domain. This option is required during creation; it may not be modified once set.

flow-eviction-policy

Specifies a flow eviction policy for the route domain to use, to select which flows to evict when the number of connections approaches the connection limit on the route domain. The default value is none.

fw-enforced-policy

Specifies an enforced firewall policy. fw-enforced-policy rules are enforced on a route-domain.

fw-enforced-policy-rules

Specifies firewall rules enforced on net route-domain via referenced fw-enforced-policy.

fw-staged-policy

Specifies a staged firewall policy. fw-staged-policy rules are not enforced while all the visibility aspects namely statistics, reporting and logging function as if the fw-staged-policy rules were enforced on a route-domain.

fw-staged-policy-rules

Specifies firewall rules staged on net route-domain via referenced fw-staged-policy.

security-nat-rules

Specifies security nat rules associated with net route-domain via referenced security-nat-policy.

parent

Specifies the route domain the system searches when it cannot find a route in the configured domain. The default value is None.

If you specify a parent, during route table lookup, if the system cannot find a route in the current route domain, the system searches routes in the parent route domain. If no route is found in the parent route domain, the system searches the parent route domain's parent, and so on, until the system finds either a match or a parent with a value of None. For example, if rd_1 has a parent of rd_0 (in this example, route domain rd_0 has a parent of None), and you include vlan_a in rd_1, when requests arrive for vlan_a, the system looks in rd_1 for a route for the specified destination. If no route is found, the system searches route domain 0. If it still cannot find a route, the request for vlan_a fails. If, using the same example, you set the parent to None, under the same conditions, the system looks in rd_1, and if it cannot find a matching route, the system refrains from searching any other route domain, the request for vlan_a fails.

port-misuse

Used to show or reset port misuse policy statistics for the route domain.

fw-context-stat

Used to show or reset firewall statistics for the route domain.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

security-nat-policy

Configures the Security NAT Policy (see security nat policy). If specified, this is the NAT policy used to perform first-match classification for incoming traffic to a virtual server if 'the virtual server itself does not have a NAT policy configured AND security-nat-policy.use-route-domain-policy is enabled on the virtual'.

service-policy

Configures the service policy for the route-domain. If set, it will enforce the service policy for incoming network traffic. The service policy can be used to set specific policy based configurations like flow timers, which applies to the flows that matches the policy specification.

strict

Specifies whether the system allows a connection to span route domains. The default value is enabled.

Note: When you enable this option, the system may find invalid iRules(r) that passed validation.

routing-protocol

Specifies routing protocols, by name, for the system to use in the route domain. The default value is none. Dynamic routing must be licensed to use this option.

vlan

Specifies VLANs, by name, for the system to use in the route domain. The default value is none.

ip-intelligence-categories

Used to show/ reset statistics on IP intelligence white/ black lists categories.

SEE ALSO

create, delete, edit, glob, list, modify, security nat policy, net service-policy, net vlan, net vlan-group, regex, show, tmsh, net bwc-policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2017-09-06 net route-domain(1)

net route

NAME

route - Configures a route for traffic management.

MODULE

net

SYNTAX

Configure the route component within the net module using the syntax in the following sections.

CREATE/MODIFY

create route [name | ip address/netmask | default | default-inet6]
modify route [name | ip address/netmask | default | default-inet6]

options:

blackhole
description [string]
gw [ip address]
interface [name]
mtu [integer]
network [ip address/netmask]
pool [name]

edit route

[[name | ip address/netmask | default | default-inet6] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list route

list route

[[name | ip address/netmask | default | default-inet6] | [glob] | [regex]] ...]

show running-config route

show running-config route

[[name | ip address/netmask | default | default-inet6] | [glob] | [regex]] ...]

options:

all-properties
mtu
non-default-properties
one-line
partition

show route

show route

[[name | ip address/netmask | default | default-inet6] | [glob] | [regex]] ...]

options:

connected
dynamic
field-fmt
lookup [ip address]
static

DELETE
delete route [name | ip address/netmask | default | default-inet6]

DESCRIPTION

You can configure routes for the system, including default routes.

Note that when you use the command edit to create a new route, by default the gw (gateway) option is set. If you do not want to use the gw option, remove that line of syntax in the editor.

EXAMPLES

```
create route myRoute3 network 12.12.4.0/24 interface external
```

Sets the route myRoute3 to the address 12.12.4.0/24 on the interface named external.

```
create route 12.12.3.0/24 gw 10.10.10.254
```

Sets the route to the subnet 12.12.3.0/24 whose gateway IP address is 10.10.10.254.

```
create route default gw 10.10.10.254
```

Sets the default gateway IP address to 10.10.10.254.

```
show route lookup myRoute
```

Displays the route that the system uses to reach the IP address 12.12.3.0.

OPTIONS

Note: The options blackhole, gw, interface, and pool are mutually exclusive. You can use only one of these options at a time, and you must specify at least one of these options when configuring a route.

blackhole

Specifies that the system drops traffic that is addressed to the specified destination.

connected

Displays connected routes.

description

User defined description.

dynamic

Displays dynamic routes.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

gw Specifies a gateway address for the system.

interface

Specifies the tunnel, VLAN or VLAN group to which the system sends traffic.

ip address/netmask

Specifies the destination subnet and mask using CIDR notation, such as 12.12.3.0/24. You can also specify the keyword default or default-inet6.

lookup

Displays the route that the system uses to reach the specified IP address. You can specify only a single IP address with the lookup option.

mtu Sets a specific maximum transition unit (MTU). If you set this option to 0 (zero), the system selects the appropriate MTU for the route, and does not display the MTUs.

network

Specifies the destination subnet and mask using CIDR notation, such as 12.12.3.0/24. You can also specify the keyword default or default-inet6.

partition

Displays the administrative partition within which the route resides.

pool Specifies a pool to which the system sends traffic. This allows the system to send traffic to multiple, load-balanced gateways.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

static

Displays static routes.

SEE ALSO

create, delete, edit, glob, list, ltm pool, modify, net vlan, net vlan-group, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

net router-advertisement

NAME

router-advertisement - Configures IPv6 prefixes for router advertisement on a VLAN.

MODULE

net

SYNTAX

Modify the router-advertisement component within the net module using the syntax shown in the following sections.

CREATE/MODIFY

create router-advertisement [name]

modify router-advertisement [name]

options:

app-service [[string] | none]

current-hop-limit [integer]

description [string]

disabled | enabled

max-interval [integer]

min-interval [integer]

mtu [integer]

no-other-config | other-config

prefixes

[add | delete | modify | replace-all-with] {

[name] ... {

app-service [[string] | none]

autonomous | not-autonomous

description [string]

on-link | not-on-link

preferred-lifetime [integer]

prefix [ip address]

prefix-length [integer]

valid-lifetime [integer]

}

}

reachable-time [integer]

retransmit-timer [integer]

router-lifetime [integer]

unmanaged | managed

vlan [name]

edit router-advertisement [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list router-advertisement

list router-advertisement [[[name] | [glob] | [regex]] ...]

show running-config router-advertisement

show running-config router-advertisement

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete router-advertisement [name]

DESCRIPTION

Router advertisements are part of the configuration of BIG-IP(r) network components. When creating a router advertisement, you must specify a VLAN on the command line.

EXAMPLES

create router-advertisement my_ra vlan my_vlan

Creates the router advertisement my_ra that includes the VLAN my_vlan.

delete router-advertisement my_ra

Deletes the router advertisement named my_ra and all associated prefixes.

OPTIONS

Note the following information regarding options for the router-advertisement component:

- Â· The options disabled and enabled are mutually exclusive.
- Â· The options no-other-config and other-config are mutually exclusive.
- Â· The options unmanaged and managed are mutually exclusive.
- Â· The options autonomous and not-autonomous are mutually exclusive.
- Â· The options on-link and not-on-link are mutually exclusive.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

autonomous

Indicates that the Autonomous Flag field in the prefix information option be set to 1. The default value is 1.

current-hop-limit

Defines the hop limit sent in the router advertisement. The default value is 0 (zero).

description

User defined description.

disabled

Disables router advertisement for the VLAN. This is the default.

enabled

Enables router advertisement for the VLAN.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

managed

Indicates that the Managed address configuration flag field in the router advertisement be set to 1.

max-interval

Specifies the maximum time allowed between sending unsolicited multicast Router Advertisements from the interface, in seconds. The default value is 600.

min-interval

Specifies, in seconds, the minimum time allowed between sending unsolicited multicast Router Advertisements from the interface. The default value is 200.

mtu Sets a specific maximum transition unit (MTU) for the VLAN. The default value is 0 (zero).

name Specifies a unique name for the component. This option is required for the create, delete, and modify commands.

no-other-config

Indicates that the Other Configuration flag field in the router advertisement be set to 0 (zero). The default value is 0 zero.

not-autonomous

Indicates that the Autonomous flag field in the prefix information option be set to 0 (zero).

not-on-link

Indicates that the on-link flag field in the prefix information option be set to 0 (zero).

on-link

Indicates that the on-link flag field in the prefix information option be set to 1. The default value is 1.

other-config

Indicates that the Other Configuration flag field in the router advertisement be set to 1.

preferred-lifetime

Specifies, in seconds, the value for the Preferred Lifetime field in the prefix information option. The default value is 604800.

prefix

Specifies the prefix for the prefix information option.

prefix-length

Specifies the length of the prefix for the prefix information option.

prefixes

Specifies the objects that hold the prefix specific information for the router advertisement.

reachable-time

Specifies the value to be used for the Reachable Time field in the Router Advertisement. The default value is 0 (zero).

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

retransmit-timer

Specifies the value to be used for the Retransmit Timer field in the Router Advertisement. The default value is 0 (zero).

router

Specifies that the router advertisement acts as a router for the VLAN.

router-lifetime

Specifies the value to be used for the Router Lifetime field in the Router Advertisement. The default value is 1800.

unmanaged

Specifies that the Managed address configuration flag field in the router advertisement be set to 0 (zero). The default value is 0 (zero).

valid-lifetime

Specifies, in seconds, the value for the Valid Lifetime field in the prefix information option. The default value is 2592000.

SEE ALSO

create, delete, edit, glob, list, modify, net vlan, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-05-22 net router-advertisement(1)

net routing access-list

NAME Early Access - access-list - placeholder

MODULE net routing

SYNTAX

CREATE/MODIFY

create access-list [name]

modify access-list [name]

options:

description [[string] | none]

route-domain [[string] | none]

entries [add | delete | modify | replace-all-with] {
[name] }

options:

action [[string] | none]

destination [ip address]

exact-match [disabled | enabled]

source [ip address]

}

}

DISPLAY

DELETE delete access-list [name]

DESCRIPTION placeholder

EXAMPLES

OPTIONS

description

route-domain

entries

action

destination

exact-match

source

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

net routing bfd

NAME Early Access - bfd - placeholder

MODULE net routing

SYNTAX

CREATE/MODIFY

create bfd [name]

modify bfd [name]

options:

gtsm [disabled | enabled]

gtsm-ttl [integer]

notification [disabled | enabled]

route-domain [[string] | none]

slow-timer [integer]

multihop-peer [add | delete | modify | replace-all-with] {

[name] }

options:

interval [[string] | none]

minrx [[string] | none]

multiplier [[string] | none]

}

vlan [add | delete | modify | replace-all-with] {

[name] }

options:

enabled [true | false]

interval [[string] | none]

minrx [[string] | none]

multiplier [[string] | none]

}

}

DISPLAY

DELETE delete bfd [name]

DESCRIPTION placeholder

EXAMPLES

OPTIONS

gtsm State of Generalized TTL Security Mechanism (GTSM) protection.

gtsm-ttl

Sets a BFD GTSM TTL (time to live) value. (default: 255)

notification

State of BFD notification.

route-domain

slow-timer

Set a BFD slow timer interval. (default: 1000)

multihop-peer

interval

minrx

multiplier

vlan

enabled

interval

minrx

multiplier

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

net routing bgp

NAME Early Access - bgp - BGP Instance

MODULE net routing

SYNTAX

CREATE/MODIFY

create bgp [name]

modify bgp [name]

options:

allow-infinite-hold-time [disabled | enabled]

always-compare-med [disabled | enabled]

bestpath {

as-path-ignore [disabled | enabled]

compare-confed-asp-path [disabled | enabled]

compare-originator-id [disabled | enabled]

compare-routerid [disabled | enabled]

med {

confed [disabled | enabled]

missing-as-worst [disabled | enabled]

remove-recv-med [disabled | enabled]

remove-send-med [disabled | enabled]

}

tie-break-on-age [disabled | enabled]

}

client-to-client-reflection [disabled | enabled]

cluster-id [integer]

confederation {

identifier [integer]

peers [[string] | none]

}

dampening {

reachability-half-life [integer]

reuse [integer]

route-map [[string] | none]

state [disabled | enabled]

suppress [integer]

suppress-max [integer]

unreachability-half-life [integer]

}

default-local-preference [integer]

description [[string] | none]

deterministic-med [disabled | enabled]

enabled [true | false]

enforce-first-as [disabled | enabled]

fast-external-failover [disabled | enabled]

graceful-restart {

graceful-reset [disabled | enabled]

restart-time [integer]

stalepath-time [integer]

}

graceful-shutdown {

capable [disabled | enabled]

local-preference [integer]

mode [disabled | enabled]

}

hold-time [integer]

keep-alive [integer]

local-as [integer]

log-neighbor-changes [disabled | enabled]

profile [[string] | none]

route-domain [[string] | none]

router-id [ip address]

scan-time [integer]

synchronization [disabled | enabled]

update-delay [integer]

view [disabled | enabled]

address-family [add | delete | modify | replace-all-with] {

[[name]] {

options:

auto-summary [disabled | enabled]

distance {

external [integer]

internal [integer]

local [integer]

}

network-synchronization [disabled | enabled]

aggregate-address [add | delete | modify | replace-all-with] {

[[name]] {

options:

as-set [disabled | enabled]

summary-only [disabled | enabled]

}

}


```

next-hop-self [disabled | enabled]
prefix-list {
  in [[string] | none]
  out [[string] | none]
}
remove-private-as [disabled | enabled]
route-map {
  in [[string] | none]
  out [[string] | none]
}
route-reflector-client [disabled | enabled]
route-server-client [disabled | enabled]
send-community [[string] | none]
soft-reconfiguration-inbound [disabled | enabled]
unsuppress-map [[string] | none]
weight [[string] | none]
}
}
}
network [add | delete | modify | replace-all-with] {
  [ [name] ] {
options:
backdoor [disabled | enabled]
route-map [[string] | none]
}
}
peer-group [add | delete | modify | replace-all-with] {
  [ [name] ] {
options:
advertisement-interval [integer]
allow-infinite-hold-time [disabled | enabled]
as-origination-interval [integer]
capability {
  dynamic [disabled | enabled]
  route-refresh [disabled | enabled]
}
capability-negotiate {
  override [disabled | enabled]
  state [disabled | enabled]
  strict-match [disabled | enabled]
}
collide-established [disabled | enabled]
connect-timer [integer]
description [[string] | none]
ebgp-multihop [integer]
enabled [true | false]
enforce-multihop [disabled | enabled]
fall-over [[string] | none]
graceful-shutdown {
  mode [disabled | enabled]
  timer [integer]
}
hold-time [integer]
keep-alive [integer]
local-as [integer]
passive [disabled | enabled]
password [[string] | none]
port [integer]
remote-as [integer]
restart-time [integer]
update-source [[string] | none]
version [integer]
address-family [add | delete | modify | replace-all-with] {
  [ [name] ] {
options:
activate [disabled | enabled]
allow-as-in [[string] | none]
as-override [disabled | enabled]
attribute-unchanged {
  as-path [disabled | enabled]
  med [disabled | enabled]
  next-hop [disabled | enabled]
}
capability {
  graceful-restart [disabled | enabled]
  orf {
prefix-list [[string] | none]
}
}
}
default-originate {
  route-map [[string] | none]
  state [disabled | enabled]
}
distribute-list {
  in [[string] | none]
  out [[string] | none]
}

```

```

}
filter-list {
  in [[string] | none]
  out [[string] | none]
}
maximum-prefix {
  threshold [[string] | none]
  value [integer]
  warning-only [disabled | enabled]
}
next-hop-self [disabled | enabled]
prefix-list {
  in [[string] | none]
  out [[string] | none]
}
remove-private-as [disabled | enabled]
route-map {
  in [[string] | none]
  out [[string] | none]
}
route-reflector-client [disabled | enabled]
route-server-client [disabled | enabled]
send-community [[string] | none]
soft-reconfiguration-inbound [disabled | enabled]
unsuppress-map [[string] | none]
weight [[string] | none]
}
}
}
}

```

DISPLAY

DELETE delete bgp [name]

DESCRIPTION None

EXAMPLES

OPTIONS

allow-infinite-hold-time

BGP timer disallow-infinite-hold-time. (default: True)

always-compare-med

option to compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. (default: False)

bestpath

as-path-ignore

prevent the router from considering the autonomous system (AS) path length as a factor in the algorithm for choosing a best path route.

compare-confed-aspath

allow comparing of the confederation AS path length.

compare-originator-id

change the default bestpath selection by not comparing an originator-ID for an identical eBGP path. (default: True)

compare-routerid

compare router IDs for identical eBGP paths.

med

confed

Compare MED among confederation paths.

missing-as-worst

Treat missing MED as the least preferred one.

remove-recv-med

To remove rcvd MED attribute.

remove-send-med

To remove send MED attribute.

tie-break-on-age

always select a preferred older route even when the bestpath.compare-routerid is enabled. (default: True)

client-to-client-reflection

Configure client-to-client route reflection. (default: True)

cluster-id

Route-reflector cluster-id in IP address or 32-bit quantity format.

confederation

identifier

Set routing domain confederation AS.

peers

AS numbers of eBGP peers that are in the same confederation.

dampening
reachability-half-life
Reachability half-life time for the penalty(minutes).

reuse
Value to start reusing a route.

route-map
Route-map to specify the criteria for dampening.

state
Enable route-flap dampening.

suppress
Value to start suppressing a route.

suppress-max
Maximum duration to suppress a stable route(minutes).

unreachability-half-life
Un-reachability half-life time for the penalty(minutes).

default-local-preference
Configure default local preference value (higher=more preferred). (default: 100)

description
deterministic-med
Pick the best-MED path among paths advertised from the neighboring AS. (default: False)

enabled
enforce-first-as
Enforce the first AS for EBGp routes.

fast-external-failover
Immediately reset session if a link to a directly connected external peer goes down. (default: True)

graceful-restart
graceful-reset
Graceful reset capability.

restart-time
Max time needed for neighbor(s) to restart (seconds).

stalepath-time
Max time to retain stale paths from restarting Neighbor(s) (seconds).

graceful-shutdown
capable
Set router as g-shut capable. (default: False)

local-preference
Set local preference to use during g-shut. (default: 0)

mode Gracefully shutdown this router. (default: False)

hold-time
Globally set or reset the holdtime value for all of the neighbors. (default: 90)

keep-alive
Globally set or reset the keepalive value for all of the neighbors. (default: 30)

local-as
log-neighbor-changes
Log neighbor up/down and reset reason. (default: False)

profile
route-domain
router-id
Manually override current router identifier (peers will reset).

scan-time
Configure background scan interval (sec). (default: 60)

synchronization
Enable IGP synchronization of internal BGP (iBGP) learned routes. (default: False)

update-delay
Max time to defer initial route-selection (sec). (default: 120)

view BGP view name.

address-family
auto-summary
enable sending summarized routes by a BGP speaker to its peers. (default: False)

distance

external
Define administrative distance for routes external to the AS. (default: 20)

internal
Define administrative distance for routes internal to the AS. (default: 200)

local
Define administrative distance for local routes. (default: 200)

network-synchronization
Enable IGP synchronization on network routes.

aggregate-address
as-set
summary-only
redistribute
route-map
 distance
 access-list
 distance
 neighbor
advertisement-interval
Minimum interval between sending BGP routing updates.

allow-infinite-hold-time
BGP per neighbor allow-infinite-hold-time. (default: True)

as-origination-interval
Minimum interval between sending AS-origination routing updates.

capability
dynamic
 Advertise dynamic capability to this neighbor. (default: False)

route-refresh
 Advertise route-refresh capability to this neighbor. (default: True)

capability-negotiate
override
 Override capability negotiation result. (default: False)

state
 Perform capability negotiation. (default: True)

strict-match
 Strict capability negotiation match. (default: False)

collide-established
Include Neighbor in Established State for Collision Detection.

connect-timer
BGP connect timer. (default: 0)

description
User defined description..

ebgp-multihop
Allow EBGp neighbors not on directly connected networks. (default: 1)

enabled
enforce-multihop
Enforce EBGp neighbors to perform multihop. (default: False)

fall-over
Fall-over detection - Bidirectional Forwarding Detection (BFD) or BFD Multihop.

graceful-shutdown
mode Gracefully shut down this neighbor. (default: False)

timer
 Max time needed for Neighbor to shutdown. (default: 60)

hold-time
Holdtime. (default: 90)

keep-alive
Keepalive interval. (default: 30)

local-as
Specify AS number to use with BGP neighbor.

passive
Don't send open messages to this neighbor. (default: False)

password
Set password to the neighbor.

peer-group
port Neighbor's BGP port. (default: 179)

remote-as
Specify AS number of BGP neighbor.

restart-time
Max time needed for Neighbor to restart. (default: 0)

update-source
Source of routing updates.

version
Set the BGP version to match a neighbor. (default: 4)

vlan
address-family
activate
Enable the address family for this neighbor. (default: True)

allow-as-in
Accept as-path with my AS present in it.

as-override
Override AS path.

attribute-unchanged
as-path
BGP attribute, as-path, is propagated unchanged to this neighbor.

med BGP attribute, med, is propagated unchanged to this neighbor.

next-hop
BGP attribute, next-hop, is propagated unchanged to this neighbor.

capability
graceful-restart
The graceful-restart capability. (default: True)

orf
prefix-list
Advertise ORF capability to peer. (default: disabled)

default-originate
route-map
Route-map to specify criteria to originate default.

state
Originate default route to this neighbor. (default: False)

distribute-list
in Filter updates to/from this neighbor.

out Filter updates to/from this neighbor.

filter-list
in Establish BGP filters.

out Establish BGP filters.

maximum-prefix
threshold
Maximum number of prefix threshold value. (default: 0)

value
Maximum number of prefix accept from this peer. (default: 0)

warning-only
Only give a warning message when the maximum-prefix.threshold limit is exceeded. (default: True)

next-hop-self
Disable the next hop calculation for this neighbor. (default: False)

prefix-list
in Filter updates to/from this neighbor.

out Filter updates to/from this neighbor.

remove-private-as
Remove private AS number from outbound updates. (default: False)

route-map
in Apply route map to neighbor.

out Apply route map to neighbor.

route-reflector-client
Configure a neighbor as Route Reflector client. (default: False)

route-server-client
Configure a neighbor as Route Server client. (default: False)

send-community
Send community attribute to this neighbor. (default: both)

soft-reconfiguration-inbound
(default: False)

unsuppress-map
Route-map to selectively unsuppress suppressed routes.

weight
Set default weight for routes from this neighbor.

 network
backdoor
route-map
 peer-group
advertisement-interval
Minimum interval between sending BGP routing updates.

 allow-infinite-hold-time
BGP per neighbor allow-infinite-hold-time. (default: True)

 as-origination-interval
Minimum interval between sending AS-origination routing updates.

 capability
dynamic
 Advertise dynamic capability to this neighbor. (default: False)

route-refresh
Advertise route-refresh capability to this neighbor. (default: True)

 capability-negotiate
override
 Override capability negotiation result. (default: False)

state
Perform capability negotiation. (default: True)

strict-match
Strict capability negotiation match. (default: False)

 collide-established
Include Neighbor in Established State for Collision Detection.

 connect-timer
BGP connect timer. (default: 0)

 description
User defined description..

 ebgp-multihop
Allow EBGp neighbors not on directly connected networks. (default: 1)

 enabled
enforce-multihop
Enforce EBGp neighbors to perform multihop. (default: False)

 fall-over
Fall-over detection - Bidirectional Forwarding Detection (BFD) or BFD Multihop.

 graceful-shutdown
mode Gracefully shut down this neighbor. (default: False)

timer
Max time needed for Neighbor to shutdown. (default: 60)

 hold-time
Holdtime. (default: 90)

 keep-alive
Keepalive interval. (default: 30)

 local-as
Specify AS number to use with BGP neighbor.

 passive
Don't send open messages to this neighbor. (default: False)

 password
Set password to the neighbor.

port Neighbor's BGP port. (default: 179)

remote-as
Specify AS number of BGP neighbor.

restart-time
Max time needed for Neighbor to restart. (default: 0)

update-source
Source of routing updates.

version
Set the BGP version to match a neighbor. (default: 4)

address-family
activate
Enable the address family for this neighbor. (default: True)

allow-as-in
Accept as-path with my AS present in it.

as-override
Override AS path.

attribute-unchanged
as-path
BGP attribute, as-path, is propagated unchanged to this neighbor.

med
BGP attribute, med, is propagated unchanged to this neighbor.

next-hop
BGP attribute, next-hop, is propagated unchanged to this neighbor.

capability
graceful-restart
The graceful-restart capability. (default: True)

orf
prefix-list
Advertise ORF capability to peer. (default: disabled)

default-originate
route-map
Route-map to specify criteria to originate default.

state
Originate default route to this neighbor. (default: False)

distribute-list
in Filter updates to/from this neighbor.

out Filter updates to/from this neighbor.

filter-list
in Establish BGP filters.

out Establish BGP filters.

maximum-prefix
threshold
Maximum number of prefix threshold value. (default: 0)

value
Maximum number of prefix accept from this peer. (default: 0)

warning-only
Only give a warning message when the maximum-prefix.threshold limit is exceeded. (default: True)

next-hop-self
Disable the next hop calculation for this neighbor. (default: False)

prefix-list
in Filter updates to/from this neighbor.

out Filter updates to/from this neighbor.

remove-private-as
Remove private AS number from outbound updates. (default: False)

route-map
in Apply route map to neighbor.

out Apply route map to neighbor.

route-reflector-client

Configure a neighbor as Route Reflector client. (default: False)

route-server-client

Configure a neighbor as Route Server client. (default: False)

send-community

Send community attribute to this neighbor. (default: both)

soft-reconfiguration-inbound

(default: False)

unsuppress-map

Route-map to selectively unsuppress suppressed routes.

weight

Set default weight for routes from this neighbor.

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

Following keywords have different identifier in legacy command line

TMOS | IMI | Note

TMOS	IMI	Note
allow-infinite-hold-time	disallow-infinite-holdtime	BGP timer disallow-infinite-hold-time.
capability-negotiate.override	override-capability	Override capability negotiation result.
capability-negotiate.state	dont-capability-negotiate	Perform capability negotiation.
capability-negotiate.strict-match	strict-capability-match	Strict capability negotiation match.
graceful-shutdown.mode	g-shut	Gracefully shut down this neighbor

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 net routing bgp(1)

net routing community-list

NAME Early Access - community-list - placeholder

MODULE net routing

SYNTAX

CREATE/MODIFY

create community-list [name]

modify community-list [name]

options:

description [[string] | none]

route-domain [[string] | none]

type [[string] | none]

entries [add | delete | modify | replace-all-with] {
[name] }

options:

action [[string] | none]

community [[string] | none]

regex [[string] | none]

}

}

DISPLAY

DELETE delete community-list [name]

DESCRIPTION placeholder

EXAMPLES

OPTIONS

description

route-domain

type

entries

action

community

regex

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 net routing community-list(1)

net routing debug

NAME Early Access - debug - placeholder

MODULE net routing

SYNTAX

CREATE/MODIFY

create debug [name]

modify debug [name]

options:

bfd [event | ipc-error | ipc-event | nsm | packet | session | all]

bgp [bfd | dampening | events | filters | fsm | keepalives | nht | nsm | updates-in | updates-out | updates | all]

nsm [events | packet | packet-send | packet-recv | packet-detail | kernel | ha | ha-all | bfd | all]

route-domain [[string] | none]

DISPLAY

DELETE delete debug [name]

DESCRIPTION placeholder

EXAMPLES

OPTIONS

bfd

bgp

nsm

route-domain

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 net routing debug(1)

net routing extcommunity-list

NAME Early Access - extcommunity-list - placeholder

MODULE net routing

SYNTAX

CREATE/MODIFY

create extcommunity-list [name]

modify extcommunity-list [name]

options:

description [[string] | none]

route-domain [[string] | none]

type [[string] | none]

entries [add | delete | modify | replace-all-with] {

[[name]] {

options:

action [[string] | none]

regex [[string] | none]

rt [[string] | none]

soo [[string] | none]

}

}

DISPLAY

DELETE delete extcommunity-list [name]

DESCRIPTION placeholder

EXAMPLES

OPTIONS

description
route-domain
type
entries
action
regex
rt
soo

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 net routing extcommunity-list(1)

net routing prefix-list

NAME Early Access - prefix-list - placeholder

MODULE net routing

SYNTAX

CREATE/MODIFY

create prefix-list [name]

modify prefix-list [name]

options:

description [[string] | none]

route-domain [[string] | none]

entries [add | delete | modify | replace-all-with] {
[name] }

options:

action [[string] | none]

prefix [ip address]

prefix-len-range [[string] | none]

}

}

DISPLAY

DELETE delete prefix-list [name]

DESCRIPTION placeholder

EXAMPLES

OPTIONS

description

route-domain

entries

action

prefix

prefix-len-range

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 net routing prefix-list(1)

net routing profile bgp

NAME Early Access - bgp - placeholder

MODULE net routing profile

SYNTAX

```

CREATE/MODIFY
create bgp [name]
modify bgp [name]
options:
  adj-out [disabled | enabled]
  aggregate-nexthop-check [disabled | enabled]
  as-local-count [integer]
  bgp-multiple-instance [disabled | enabled]
  defaults-from [[string] | none]
  description [[string] | none]
  extended-asn-cap [disabled | enabled]
  max-paths {
  ebgp [integer]
  ibgp [integer]
  }
  nexthop-trigger {
  delay [integer]
  state [disabled | enabled]
  }
  rfc1771 {
  path-select [disabled | enabled]
  strict [disabled | enabled]
  }
  router-id [ip address]

```

```

DISPLAY
DELETE delete bgp [name]
DESCRIPTION placeholder
EXAMPLES
OPTIONS

```

```

  adj-out
  Disable BGP ADJ_OUT. (default: True)

```

```

  aggregate-nexthop-check
  Option to perform aggregation only when next-hop matches the specified IP address..

```

```

  as-local-count
  Set the number of times the local-AS (Autonomous System) is to be prepended. (default: 1)

```

```

  bgp-multiple-instance
  Enable to allow BGP views and disallow BGP graceful-restart (peers will reset). (default: False)

```

```

  defaults-from
  description
  extended-asn-cap
  Enable the router to send 4-octet ASN capabilities. (default: True)

```

```

  max-paths
  ebgp Set number of equal-cost multi-path (ECMP) routes for eBGP.

```

```

  ibgp Set number of equal-cost multi-path (ECMP) routes for iBGP.

```

```

  nexthop-trigger
  delay
  Configure nexthop trigger delay time interval. (default: 5)

```

```

  state
  Enable the nexthop tracking functionality. (default: False)

```

```

  rfc1771
  path-select
  Set RFC1771 compatible path selection mechanism. (default: False)

```

```

  strict
  Set the origin path attribute to "IGP" when the origin is a protocol such as RIP, OSPF, or ISIS.
  (default: False)

```

```

  router-id
  Manually override current router identifier (peers will reset).

```

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

Following keywords have different identifier in legacy command line

TMOS	IMI	Note
adj-out	disable-adj-out	Disable BGP ADJ_OUT.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

net routing route-map

NAME Early Access - route-map - placeholder

MODULE net routing

SYNTAX

CREATE/MODIFY

create route-map [name]

modify route-map [name]

options:

description [[string] | none]

route-domain [[string] | none]

entries [add | delete | modify | replace-all-with] {
[name] }

options:

action [[string] | none]

match {

as-path [[string] | none]

community {

exact-match [[string] | none]

list [[string] | none]

}

extcommunity {

exact-match [[string] | none]

list [[string] | none]

}

ipv4 {

address {

access-list [[string] | none]

prefix-list [[string] | none]

}

next-hop {

access-list [[string] | none]

prefix-list [[string] | none]

}

peer {

access-list [[string] | none]

}

}

ipv6 {

address {

access-list [[string] | none]

prefix-list [[string] | none]

}

next-hop {

access-list [[string] | none]

prefix-list [[string] | none]

}

peer {

access-list [[string] | none]

}

}

metric [integer]

origin [[string] | none]

route-type [[string] | none]

tag [integer]

vlan [[string] | none]

}

set {

aggregator {

address [ip address]

as [integer]

}

as-path-prepend [integer]

atomic-aggregate [[string] | none]

community {

additive [[string] | none]

exact-set [[string] | none]

value [integer]

}

dampening {

reachability-half-life [integer]

reuse [integer]

suppress [integer]

suppress-max [integer]

unreachability-half-life [integer]

}

extcommunity {

```

rt [[string] | none]
soo [[string] | none]
}
ip {
next-hop {
  address [ip address]
}
}
ipv6 {
next-hop {
  address [ip address]
  local [ip address]
}
}
level [[string] | none]
local-preference [integer]
metric {
type [[string] | none]
value [[string] | none]
}
origin [[string] | none]
originator-id [ip address]
tag [integer]
weight [integer]
}
}
}

```

DISPLAY

DELETE delete route-map [name]

DESCRIPTION placeholder

EXAMPLES

OPTIONS

```

description
route-domain
entries
action
match
as-path
  Define a BGP AS path access list.

```

community

```

exact-match
Do exact matching of communities. (default: False)

```

```

list Match BGP ecommunity list.

```

extcommunity

```

exact-match
Do exact matching of extcommunities. (default: False)

```

```

list Match BGP ecommunity list.

```

ipv4

```

address
access-list
Match entries of IPv4 access-lists.

```

```

prefix-list
Match entries of IPv4 prefix-lists.

```

```

next-hop
access-list
Match entries of IPv4 access-lists.

```

```

prefix-list
Match entries of IPv4 prefix-lists.

```

```

peer
access-list
Match entries of IPv4 access-lists.

```

ipv6

```

address
access-list
Match entries of IPv6 access-lists.

```

```

prefix-list
Match entries of IPv6 prefix-lists.

```

```

next-hop
access-list
Match entries of IPv6 access-lists.

```

```

prefix-list
Match entries of IPv6 prefix-lists.

```

peer
access-list
Match entries of IPv6 access-lists.

metric
Match metric of route.

origin
Match BGP origin code.

route-type
Match route type.

tag Match tag.

vlan Match first hop interface of route.

set
aggregator
address
set BGP aggregator IP address.

as set BGP aggregator attribute.

as-path-prepend
Prepend string for a BGP AS-path attribute.

atomic-aggregate
BGP atomic aggregate attribute.

community
additive
Add to the existing community.

exact-set
Do exact setting of communities. (default: False)

value
set community number, either in AA:NNO format, or number..

dampening
reachability-half-life
Reachability half-life time for the penalty(minutes). (default: 15)

reuse
Value to start reusing a route. (default: 750)

suppress
Value to start suppressing a route. (default: 2000)

suppress-max
Maximum duration to suppress a stable route(minutes). (default: 60)

unreachability-half-life
Un-reachability half-life time for the penalty(minutes). (default: 1)

extcommunity
rt VPN extended community.

soo Site-of-Origin extended community.

ip
next-hop
address
IP address of next-hop.

ipv6
next-hop
address
IPv6 local address.

local
IPv6 local address.

level
IP address of next-hop.

local-preference
BGP local preference path attribute.

metric
type Metric value.

value

origin
set BGP origin code.

originator-id
set IP address of originator.

tag set tag.

weight
set BGP weight for routing table.

SEE ALSO

NOTES

This is an early access feature, experimental and susceptible to change in future releases.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 net routing route-map(1)

net rst-cause

NAME

rst-cause - Displays and Reset TCP/IP Reset Cause Statistics

MODULE

net

SYNTAX

Display and Reset the rst-cause component within the net module using the syntax in the following section.

MODIFY

reset-stats rst-cause

DISPLAY

show rst-cause

options:

(default | field-fmt)

DESCRIPTION

You can use the rst-cause component to display and reset TCP/IP reset cause statistics. This will help to debug the reason for TCP/IP reset.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011. All rights reserved.

BIG-IP 2011-01-24 net rst-cause(1)

net self-allow

NAME

self-allow - Configures the default "allow list" for all self IP addresses on the BIG-IP(r) system when the option allow-service of the component self is set to default.

MODULE

net

SYNTAX

Modify the self-allow component within the net module using the syntax shown in the following sections.

MODIFY

```
modify self-allow
options:
  defaults [all | none]
  defaults
    [add | delete | replace-all-with] {
      [protocol:port] ...
    }
}
```

```
edit self-allow
options:
  all-properties
```

```
DISPLAY
list self-allow
show running-config self-allow
options:
  all-properties
  defaults
  one-line
```

```
DELETE
You cannot delete the default allow list.
```

DESCRIPTION

You can use the self-allow component to modify or display the default allow list for all self IP addresses on the BIG-IP system when the option allow-service of the component self is set to default. The default allow list displays which service and protocol ports allow connections from outside the system. The system refuses connections made to a service or protocol port that is not on the list.

EXAMPLES

```
modify self-allow defaults all
```

Sets the default allow list to all. Then, if the value of the option allow-service of the net self component is default, the system accepts traffic from all protocol port combinations.

```
modify self-allow default replace-all-with { tcp:55 }
```

Sets the default "allow list" for all self IP addresses on the system to TCP on port 55.

```
list self-allow defaults
```

Displays the default "allow list" for all self IP addresses on the system.

OPTIONS

```
defaults
Specifies to set the default allow list to one of the following:
```

all Specifies that all protocols and services allow connections from outside the system. Use this option to open the system to complete access.

none Specifies that no protocols or services allow connections from outside the system.

```
protocol:port
Specifies a list of protocols/services that allow connections from outside the system.
```

```
replace-all-with
Specifies to replace the current protocols and services that allow connections from outside the system with the specified protocols and services.
```

SEE ALSO

edit, list, modify, net vlan, net vlan-group, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013. All rights reserved.

BIG-IP 2013-04-12 net self-allow(1)

net self

NAME

self - Configures a self IP address for a VLAN.

MODULE

net

SYNTAX

Modify the self component within the net module using the syntax shown in the following sections.

CREATE/MODIFY

```
create self [name]
modify self [name]
options:
  address [ip address/netmask]
  address-source [from-management | from-user]
  allow-service [all | default | none]
  allow-service
    [add | delete | replace-all-with] {
      [protocol:port] ...
    }
  app-service [[string] | none]
  description [string]
  fw-enforced-policy [ [policy_name] | none ]
  fw-staged-policy [ [policy_name] | none ]
  service-policy [ [policy_name] | none ]
  traffic-group [[string] | default | non-default | none]
  vlan [name]
```

```
edit self [
  [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
reset-stats self [ [ [name] | [glob] | [regex] ] ... ]
fw-enforced-policy-rules { [rule name] }
fw-staged-policy-rules { [rule name] }
options:
  fw-context-stat
```

```
mv self [ [[source-name] [destination-name]] | [[name] to-folder [folder-name]] | [[name...name] to-folder [folder-name]] ]
options:
  to-folder
```

DISPLAY

```
list self
list self
  [ [ [name] | [glob] | [regex] ] ... ]
show running-config self
show running-config self
  [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

```
show self [name]
options:
  fw-context-stat
```

DELETE

```
delete self [name]
```

DESCRIPTION

A self IP address is an IP address that is assigned to the system. Self IP addresses are part of the configuration of the BIG-IP(r) network components. You must define at least one self IP address for each VLAN.

EXAMPLES

```
create self mySelf address 10.10.10.24/16 vlan internal
```

Adds the self IP address 10.10.10.24 to the VLAN named internal. This entry is named mySelf. Alternatively, the name can encompass the IP address and mask fields, like the following example.

```
create self 10.10.10.24/16 vlan internal
```

Adds the self IP address 10.10.10.24 to the VLAN named internal.

```
modify self 10.1.1.1/16 vlan external traffic-group /Common/traffic-group-1
```

Enables a floating IP address on the external VLAN. The traffic-group option makes this virtual address available to whichever device is active on the given traffic-group. In other words, when the standby device becomes the active device for that traffic-group, it uses this virtual address. Only one of the devices in the traffic-group can use the IP address at any given time.

```
mv /net self /Common/10.10.10.15/24 /Common/myselfIP2
```

Moves/Renames the Self IP from 10.10.10.15/24 to myselfIP2.

Note: If you wish to change the name of the self IP, you may use a name that is the same as the IP Address or a name that does not represent a different IP Address than the one configured. If using prefix-length adornment on the name, it must match the existing prefix-length/netmask for the self IP.

Please refer to the mv manual page for additional examples on how to use the mv command.

Options

allow-service

Specifies the type of protocol/service that the VLAN handles. If you use this property to allow SSH, HTTP, and/or HTTPS service, administrators can use this self-IP address to log into the BIG-IP system; this makes the current self-IP available as a management-IP address on the VLAN.

The options are:

add Adds the specified protocol/service to the VLAN.

all Specifies that the VLAN handles all protocols/services.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

default

Specifies that the system uses a pre-defined set of network protocols/services that are commonly required for BIG-IP deployment. You can customize this set of services with the self-allow component.

This is not the default for the allow-service property; none, described below, is the actual default.

delete

Removes the specified protocol/service from the VLAN.

none Specifies that the VLAN handles no protocols/services. This is the default setting for a self IP address.

replace-all-with

Replaces the current protocol/service that the VLAN handles with the specified protocol/service.

address

Specifies the IP address and netmask to be assigned to the system. This is an optional field. If not specified, the name of the entry must appear in the format [ip address/mask].

address-source

Specifies the source of the self IP. This is an optional field. If not specified, the default value of from-user is used.

The options are:

from-management

Assigns the self IP with the management IP rather than the provided address or entry name.

from-user

Assigns the self IP with the provided address or entry name.

fw-context-stat

Used to show or reset firewall statistics for the self IP.

description

User-defined description.

floating

Read-only property based on the traffic-group. A floating self IP address is a self IP address for a VLAN that serves as a shared address by all devices of a BIG-IP traffic-group.

fw-enforced-policy

Specifies an enforced firewall policy. fw-enforced-policy rules are enforced on a self IP address.

fw-enforced-policy-rules

Specifies firewall rules enforced on net self via referenced fw-enforced-policy.

fw-staged-policy

Specifies a staged firewall policy. fw-staged-policy rules are not enforced while all the visibility aspects namely statistics, reporting and logging function as if the fw-staged-policy rules were enforced on a self IP address.

service-policy

Configures the service policy for the self IP address. If set, it will enforce the service policy for incoming network traffic. The service policy can be used to set specific policy based configurations like flow timers, which applies to the flows that matches the policy specification.

fw-staged-policy-rules

Specifies firewall rules staged on net self via referenced fw-staged-policy.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for

a description of regular expression syntax.

unit Read-only property that specifies the unit in a redundant system. Based on traffic-group.

traffic-group

Specifies the traffic group of the self IP address. The default traffic group is traffic-group-local-only, the non-floating traffic-group.

inherited-traffic-group

Read-only property that indicates if the traffic-group is inherited from the parent folder.

vlan Specifies the VLAN for which you are setting a self IP address. This option is required.

SEE ALSO

create, delete, edit, glob, list, modify, mv, net self-allow, net service-policy, net vlan, net vlan-group, regex, security log profile, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2017-09-06 net self(1)

net service-policy

NAME

service-policy - Configures the service policy.

MODULE

net

SYNTAX

Modify the service policy component within the net module using the syntax shown in the following sections.

CREATE/MODIFY

create service-policy [name]

modify service-policy [name]

options:

description [string]

port-misuse-policy [[port misuse policy name] | [none]]

timer-policy [[timer policy name] | [none]]

edit service-policy [[name] | all]

options:

all-properties

non-default-properties

DISPLAY

list service-policy

show running-config service-policy

options:

all-properties

non-default-properties

one-line

DESCRIPTION

service-policy configuration allows one to specify certain properties that would apply to a flow. Service policy consists of other policy objects like timer policy and port misuse policy objects. The policy can be applied at different contexts, like Global context, virtual server context, route domain context, self-ip context, or a firewall rule. When a service policy is configured both at a context level, as well as on a firewall rule, and a flow matches the rule, the more specific service-policy configuration in the rule will override the service policy setting at the context level. You can use the service-policy component to configure a shareable and reusable set of network service policies which can be associated with a number of configuration objects of the following types: net self, net route-domain, security firewall policy rules, security firewall rule-list rules, ltm virtual. Port misuse policy object is not effective with net self context.

EXAMPLES

```
create net service-policy flow-policy timer-policy idle-flow-policy
```

```
list service-policy
```

```
net service-policy flow-policy {
```

```
timer-policy idle-flow-policy
```

```
}
```

Creates service policy and associates a timer policy configuration object. (see "net timer-policy").

modify net service-policy flow-policy port-misuse-policy tcp-port-policy

```
list service-policy
net service-policy flow-policy {
  timer-policy idle-flow-policy
  port-misuse-policy tcp-port-policy
}
```

Associates a port misuse policy with service policy. (see "security firewall port-misuse-policy").

list service-policy

Displays the current service policy configuration.

OPTIONS

description
User defined description.

port-misuse-policy
Specify service port misuse policy.

timer-policy
Specify service timer policy.

SEE ALSO

create, edit, list, modify, security firewall rule-list, security firewall policy, net timer-policy, security firewall port-misuse-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-03-14 net service-policy(1)

net sfc-stats

NAME

sfc-stats - Display and Reset SFC Statistics.

MODULE

net

SYNTAX

Display and Reset the sfc-stats component within the net module using the syntax in the following section.

MODIFY

reset-stats sfc-stats

DISPLAY

show sfc-stats
options:
(default | field-fmt)

DESCRIPTION

You can use the sfc-stats component to display and reset SFC statistics. This will help to debug the SFC feature.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-09 net sfc-stats(1)

net sfc chain

NAME
chain - Defines a service function chain in the Service Function Chaining (SFC) architecture.

MODULE
net sfc

SYNTAX
Define a SFC chain.

CREATE
create chain [name]
options:
 app-service [[string] | none]
 description [[string] | none]
 hops
 [add | delete] {
[name] ... {
 app-service [[string] | none]
 description [[string] | none]
 hopkey [service-index | interface]
 [nexthop-endpoint-ip [ip-address] | nexthop-service [string] | nexthop-terminate]
 service-index [integer]
 source-interface [[string] | none]
 }
 path-id [integer]

DISPLAY
list chain
list chain [[[name] | [glob] | [regex]] ...]
show running-config net sfc chain
show running-config net sfc chain [[[name] | [glob] | [regex]] ...]
options:
 all-properties
 non-default-properties
 one-line

DELETE
delete net sfc chain [all | [name]]

DESCRIPTION
You can use the net sfc chain to create a n ordered set of abstract service functions and ordering constraints that must be applied to packets and/or frames and/or flows selected as the result of classification.

There is no modify for net sfc chain. Any modification needs to be done by deleting the chain and then creating it.

EXAMPLES
create net sfc chain SC

Creates a service function chain named SC. (See below).

list net sfc chain SC all-properties

Displays all of the properties of a service function chain.

delete net sfc chain

Deletes a service function chain (see below).

Example :

```
list net sfc chain chain1 net sfc chain chain1 {  
hops {  
  1 {  
  nexthop-service sf2-1  
  service-index 255  
  }  
  2 {  
  nexthop-terminate  
  service-index 254  
  }  
}  
path-id 678 }
```

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

name Specifies a unique name for the chain. This option is required for the commands create and delete.

description
User defined description.

hops The list of next-hops.

name Specifies a unique name for the hop. This option is required for the commands create and delete.

hopkey Specifies the type of hop, service-index or interface. The default is service-index.

nexthop-endpoint-ip
If nexthop-endpoint-ip is specified then it is assumed the next hop is NSH enabled. Only one of nexthop-endpoint-ip, nexthop-service, nexthop-terminate may be specified.

nexthop-service
The service function (SF) name of the next hop. Only one of nexthop-endpoint-ip, nexthop-service, nexthop-terminate may be specified.

nexthop-terminate
If nexthop-terminate is specified then the chain is terminated and the NSH header is removed. Only one of nexthop-endpoint-ip, nexthop-service, nexthop-terminate may be specified.

service-index
The service-index is required when hopkey is service-index.

source-interface
Specifies the interface name, a tunnel or vlan, to match on ingress for this hop. The source-interface is required when hopkey is interface.

path-id
The service path identifier that uniquely identifies the service function path.

SEE ALSO
create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-07-24 net sfc chain(1)

net sfc hop

NAME
hop - A Service Function Chain (SFC) hop object for showing statistics.

MODULE
net sfc

SYNTAX
Display the statistics for the net sfc hop objects.

DISPLAY
show running-config net sfc hop
show running-config net sfc hop [[name] | [glob] | [regex]] ...]

DESCRIPTION
Display the statistics for the net sfc hop objects.

EXAMPLE
show net sfc hop

```
-----  
Net: SFC Hop Statistics  
Hop Name Chain Name Fwd Packets In Fwd Packets Out Fwd Bytes In ...  
-----  
2 chain1 0 0 0  
1 chain1 0 0 0
```

SEE ALSO
glob, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

net sfc sf

NAME

sf - Service Function (SF)

MODULE

net sfc

SYNTAX

Configure a sf definition within the net sfc using the syntax in the following sections.

CREATE/MODIFY

create sf [name]

modify sf [name]

options:

app-service [[string] | none]

description [string]

egress-interface [[string] | none]

ingress-interface [[string] | none]

ip-address [[ip address] | none]

nsh-aware [disabled | enabled]

pool-name [[pool_name] | none]

virtual-name [[virtual_name] | none]

edit sf [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list sf

list sf [[[name] | [glob] | [regex]] ...]

show running-config net sfc sf

show running-config net sfc sf [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete net sfc sf [all | [name]]

DESCRIPTION

You can use the net sfc sf to define a service function (SF) to be referenced by a net sfc chain.

EXAMPLES

```
create net sfc sf
```

Creates a sfc sf (see below).

```
list net sfc sf all-properties
```

Displays all of the properties of all of the sfc sf.

```
delete net sfc sf
```

Deletes a sfc sf (see below).

Example:

```
net sfc sf sf2-1 {  
ip-address any  
virtual-name __f5_sfc_app__ }
```

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

name Specifies a unique name for the sf. This option is required for the commands create, delete, and modify.

description

User defined description.

egress-interface
Egress interface.

ingress-interface
Ingress interface.

ip-address
IP address of the SF. Only one of ip-address, pool-name, virtual-name may be specified.

nsh-aware
Indicate if the SF is capable of network service header (NSH) packets.

pool-name
Only one of ip-address, pool-name, virtual-name may be specified. If specified, both ingress and egress interface, vlan name, need to be specified and of the same type. If nsh-aware is enabled, then ingress and egress VXLAN-GPE tunnel names can be the same or different.

virtual-name
Only one of ip-address, pool-name, virtual-name may be specified. If specified, then no ingress or egress interface should be specified. nsh-aware cannot be specified with virtual-name defined.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tms

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

Fpp/tms/man/net_sf_sf.pod5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-07-20 net sf sf(1)

net stp-globals

NAME

stp-globals - Configures spanning tree protocols on the system.

MODULE

net

SYNTAX

Configure the stp-globals component within the net module using the syntax shown in the following sections.

MODIFY

modify stp-globals

options:

config-name [configuration name]
config-revision [integer]
description [string]
fwd-delay [integer]
hello-time [integer]
max-age [integer]
max-hops [integer]
mode [disabled | mstp | passthru | rstp | stp]
transmit-hold [integer]

edit stp-globals

options:

all-properties
non-default-properties

DISPLAY

list stp-globals

show running-config stp-globals

options:

all-properties
non-default-properties
one-line

DESCRIPTION

Provides the ability to configure spanning tree protocols for the traffic management system. Spanning tree protocols are Layer 2 protocols for preventing bridging loops. The system supports multiple spanning tree protocol (MSTP), rapid spanning tree protocol (RSTP), and spanning tree protocol (STP).

EXAMPLES

modify stp-globals mode passthru

Sets the STP mode to passthru. Passthru mode forwards spanning tree bridge protocol data units (BPDUs)

received on any interface to all other interfaces.

modify stp-globals mode disabled

Sets the STP mode to disabled. No MSTP, RSTP, or STP packets are transmitted or received on the interface or trunk, and the spanning tree algorithm exerts no control over forwarding or learning on the port or the trunk.

OPTIONS

config-name

Specifies the configuration name (1 - 32 characters in length) only when the spanning tree mode is MSTP. The default configuration name is a string representation of a globally unique MAC address belonging to the traffic management system.

The MSTP standard introduces the concept of spanning tree regions, which are groups of adjacent bridges with identical configuration names, configuration revision levels, and assignments of VLANs to spanning tree instances.

Note: The system default configuration name is a string representation of the globally unique MAC address of the traffic management system in which hyphens replace the colons in the standard MAC address. For example, the default configuration name 00-01-D7-68-11-80, represents the MAC address 00:01:D7:68:11:80.

config-revision

Specifies the revision level of the MSTP configuration only when the value of the mode option is mstp. The specified number must be in the range 0 through 65535. The default value is 0 (zero).

description

User defined description.

fwd-delay

In the original STP, the forward delay parameter controlled the number of seconds for which an interface was blocked from forwarding network traffic after a reconfiguration of the spanning tree topology. This parameter has no effect when RSTP or MSTP are used, as long as all bridges in the spanning tree use the RSTP or MSTP protocol. If any legacy STP bridges are present, then neighboring bridges must fall back to the old protocol, whose reconfiguration time is affected by the value of the fwd-delay option. The default value is 15 seconds, and the valid range is 4 to 30.

hello-time

Specifies the time interval in seconds between the periodic transmissions that communicate spanning tree information to the adjacent bridges in the network. The default value is 2 seconds, and the valid range is 1 - 10. The default value is optimal in virtually all cases. F5 Networks recommends that you do not change the value of the hello-time option.

max-age

Specifies the number of seconds for which spanning tree information received from other bridges is considered valid. The default value is 20 seconds, and the valid range is 6-40 seconds.

max-hops

Specifies the maximum number of hops an MSTP packet can travel before it is discarded. Use this option only when the value of the mode option is mstp. The number of hops must be in the range of 1 to 255 hops. The default number of hops is 20.

mode Specifies one of three spanning tree modes:

disabled

Specifies to discard spanning tree bridge protocol data units (BPDUs) received on any interface.

mstp Specifies multiple spanning tree protocol.

passthru

Forwards spanning tree bridge protocol data units (BPDUs) received on any interface to all other interfaces. Essentially, passthru mode makes the traffic management system transparent to spanning tree BPDUs. This is the system default.

rstp Specifies rapid spanning tree protocol (RSTP) converges to a fully-connected state quickly.

stp The system supports STP mode for legacy systems. If STP is detected in the network, the traffic management system changes to STP mode even when the mode option is set to disabled, mstp, or rstp.

transmit hold

Specifies the absolute limit on the number of spanning tree protocol packets the traffic management system may transmit on a port in any hello-time interval. It is used to ensure that spanning tree packets do not unduly load the network even in unstable situations. The default value is 6 packets, and the valid range is 1 through 10 packets.

SEE ALSO

edit, interface, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2016. All rights reserved.

net stp

NAME

stp - Configures a Spanning Tree Protocol (STP) instance.

MODULE

net

SYNTAX

Configure the stp component within the net module using the syntax shown in the following sections.

CREATE/MODIFY

```
create stp [all | [name] ]
modify stp [all | [name] ]
options:
  app-service [[string] | none]
  description [string]
  instance-id [integer]
  interfaces [ add | delete | modify | replace-all-with ] {
    [interface name] {
      options:
        app-service [[string] | none]
        external-path-cost [integer]
        internal-path-cost [integer]
        priority [integer]
      }
    }
  interfaces none
  priority [integer]
  trunks [ add | delete | modify | replace-all-with ] {
    [interface name] {
      options:
        app-service [[string] | none]
        external-path-cost [integer]
        internal-path-cost [integer]
        priority [integer]
      }
    }
  trunks none
  vlans [ add | delete | replace-all-with ] {
    [vlan name ...]
  }
  vlans none
```

```
edit stp [ [ all | [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
non-default-properties
```

DISPLAY

```
list stp
list stp [ [ all | [name] | [glob] | [regex] ] ... ]
show stp running-config
show stp running-config [ [ all | [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

DELETE

```
delete stp [all | [name] ]
```

DESCRIPTION

You can use the stp component to configure an STP instance.

EXAMPLES

```
list stp
```

Displays all STP instances on the system.

```
show running-config stp
```

Displays the running configuration information for all STP instances.

```
delete stp myStp2
```

Removes all members from the STP instance, and then deletes the instance itself.

Note that you cannot delete spanning tree instance 0 (the Common and Internal Spanning Tree). You can

only use the command delete in Multiple Spanning Tree Protocol (MSTP) mode.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

external-path-cost

Specifies the external path cost number for either an interface or trunk. The default value is 20000.

Each network interface has an associated path cost within each spanning tree instance. The path cost represents the relative cost of sending network traffic through that interface. In calculating the spanning tree, the algorithm tries to minimize the total path cost between each point of the tree and the root bridge. By manipulating the path costs of different interfaces or trunks it is possible to steer traffic toward paths that are faster, more reliable, and/or more economical. Path costs can take values in the range 1 - 200,000,000. The default path cost for an interface or a trunk is based on the maximum, not actual speed, of the interface or trunk.

In MSTP mode there are two kinds of path cost: external and internal. The external path cost applies only to spanning tree instance 0, the Common and Internal Spanning Tree (CIST). It is used to calculate the cost to reach an adjacent spanning tree region.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

instance-id

The instance ID for this STP instance. In Multiple Spanning Tree Protocol (MSTP) mode, there will be exactly one STP instance with ID 0. The instance ID can be a value between 1 and 255.

internal-path-cost

Specifies the internal path cost number for either an interface or trunk. The default value is 20000.

Each network interface has an associated path cost within each spanning tree instance. The path cost represents the relative cost of sending network traffic through that interface. In calculating the spanning tree, the algorithm tries to minimize the total path cost between each point of the tree and the root bridge. By manipulating the path costs of different interfaces or trunks it is possible to steer traffic toward paths that are faster, more reliable, and/or more economical. Path costs can take values in the range 1 - 200,000,000. The default path cost for an interface or a trunk is based on the maximum, not actual speed, of the interface or trunk.

In MSTP mode there are two kinds of path cost: external and internal. The internal path costs can be independently set for each spanning tree instance (including instance 0) in MSTP mode. The internal path costs are used to calculate the costs of reaching adjacent bridges within the same spanning tree region.

priority

Specifies the priority number of either a bridge, interface, or trunk. The default value for a bridge is 61440. The default value for both interfaces and trunks is 128.

Each bridge, interface, and trunk in a spanning tree instance has a priority value. The relative values of the priorities control the topology of the spanning tree chosen by the protocol. The bridge with the lowest priority value (numerically) will become the root of the spanning tree. Priority values vary from 0 - 61440 in steps of 4096.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

vlan

Specifies the VLANs that you want to add to, delete from, or replace-all-with for this STP instance.

SEE ALSO

create, delete, edit, glob, list, modify, net interface, net trunk, net vlan, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013. All rights reserved.

BIG-IP 2014-05-16 net stp(1)

NAME

timer-policy - Configures the timer policy.

MODULE

net

SYNTAX

Modify the timer policy component within the net module using the syntax shown in the following sections. A timer-policy is attached to a service-policy and applied either through an ACL rule or policy applied on a context. The list of supported contexts where a timer policy can be applied are: Virtual Server, SelfIP, Route Domain and Global. The precedence of the timer policy is as follows (highest precedence is 1):

1. ACL rule configured on a Virtual Server or SelfIP
2. Policy configured on a Virtual Server or SelfIP
3. ACL rule configured on a Route Domain
4. Policy configured on a Route Domain
5. ACL rule configured through Global Rules
6. Policy configured on Global Service Policy

Note that within the same context, ACL rule based service policy takes first precedence. Among the different contexts, the order of precedence is as follows: Virtual Server, SelfIP, Route Domain, Global.

CREATE/MODIFY

create timer-policy [name]

modify timer-policy [name]

options:

description [string]

rules [add | delete | modify | replace-all-with] {
[rule name] }

options:

description [string]

destination-ports [add | delete | replace-all-with] {
[port] | [port1-port2] }

}

destination-ports none

ip-protocol [protocol name]

timers [add | delete | modify | replace-all-with] {
[flow timer type] }

value [timeout]

}

}

timers none

}

}

rules none

edit timer-policy [[name] | all]

options:

all-properties

non-default-properties

DISPLAY

list timer-policy

show running-config timer-policy

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the timer-policy component to configure a shareable and reusable set of network timer policies which can be associated with a service policy object.

EXAMPLES

```
create net timer-policy add idle-flow-policy { rules add { r1 { ip-protocol tcp destination-ports add { 80  
8080 } timers add { flow-idle-timeout { value 120 } } } r2 { ip-protocol udp destination-ports add { 7878 }  
timers add { flow-idle-timeout { value 300 } } } }
```

```
list timer-policy
```

```
net timer-policy {
```

```
idle-flow-policy {
```

```
rules {
```

```
r1 {
```

```
ip-protocol tcp
```

```
destination-ports {
```

```
http { }
```

```
webcache { }
```

```
}
```

```
timers {
```

```
flow-idle-timeout {
```

```
value 120
```

```
}
```

```
}
```

```
}
```

```
r2 {
```

```
ip-protocol udp
```

```
destination-ports {
```


list timer-policy

Displays the current timer policy configuration list.

OPTIONS

description

User defined description.

rules

Adds, deletes, or replaces a named timer policy rule.

ip-protocol

Specifies the IP protocol entry for which the timer policy rule is being configured. This could be a layer-4 protocol (such as tcp, udp or sctp). Only flows matching the configured protocol will make use of this rule. Press the `?` key for a full list of valid protocols. Keyword 'all-other' as an ip-protocol entry means, if there are no specific ip-protocol rule that matches the flow, the flow then matches the 'all-other' ip-protocol rule. Please see example above for rule match behavior.

destination-ports

Specifies the destination port or port range to match against the flow. Keyword 'all-other' as a port entry means, if there are no specific port entry rules to match against the flow, the flow then matches the 'all-other' port rule. For eg. if a policy consists of just two rules r1 and r2, with the same protocol 'tcp' but destination port 80 for r1 and port 'all-other' for r2 configured, an incoming flow with port 80 will match r1 and incoming flow with port 9090 will match r2. Without the 'all-other' port rule r2, incoming flow with port 9090 will not match any rule.

timers

Specifies the flow timer configuration for the different timer types.

value

Specifies the timeout value in seconds.

SEE ALSO

create, edit, list, modify, security firewall rule-list, security firewall policy, net service-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-03-14 net timer-policy(1)

net trunk

NAME

trunk - Configures a trunk with link aggregation.

MODULE

net

SYNTAX

Modify the trunk component within the net module using the syntax shown in the following sections.

CREATE/MODIFY

create trunk [name]

modify trunk [name]

options:

app-service [[string] | none]

bandwidth

description [string]

distribution-hash [dst-mac | src-dst-ipport | src-dst-mac]

interfaces

[add | delete | replace-all-with] {

[name ...]

}

lACP [disabled | enabled]

lACP-mode [active | passive]

lACP-timeout [short | long]

link-select-policy [auto | maximum-bandwidth]

mac-address [MAC address]

stp [disabled | enabled]

stp-reset

qinq-ethertype [string]

edit trunk [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

reset-stats trunk

reset-stats trunk [[[name] | [glob] | [regex]] ...]

DISPLAY

list trunk

list trunk [[[name] | [glob] | [regex]] ...]

show running-config trunk

show running-config trunk

[[[name] | [glob] | [regex]] ...]

options:

all-properties

cfg-mbr-count

non-default-properties

one-line

working-mbr-count

show trunk

show trunk [[[name] | [glob] | [regex]] ...]

options:

all-properties

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

detail

field-fmt

DELETE

delete trunk [all | [name]]

DESCRIPTION

Link Aggregation allows multiple physical links to be treated as one logical link. It is also referred to as trunking.

The main objective of link aggregation is to provide increased bandwidth at a lower cost, without having to upgrade hardware. The bandwidth of the aggregated trunk is the sum of the capacity of individual member links. Thus, it provides an option for linearly incremental bandwidth as opposed to bandwidth options available through physical layer technology. The traffic management system supports link aggregation control protocol (LACP).

When a trunk is created, LACP is disabled by default. In this mode, no control packets are exchanged and the member links carry traffic as long as the physical layer is operational. In the event of physical link failure, an LACP member is removed from the aggregation.

Note that both endpoints of the trunk should have identical LACP configuration to work properly. A mixed configuration where one endpoint is LACP enabled and the other is LACP disabled, is not valid.

EXAMPLES

```
create trunk my_trunk interfaces add {1.1 1.2 1.3}
```

Creates a trunk named my_trunk that includes the interfaces 1.1, 1.2, and 1.3.

```
modify trunk my_trunk lacp enabled
```

Enable LACP on the trunk named my_trunk.

```
modify trunk my_trunk lacp-mode active
```

Enable active LACP mode on the trunk my_trunk.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

bandwidth

Specifies the operation bandwidth in bytes per second.

cfg-mbr-count

Displays the number of configured members that are associated with this trunk.

description

User defined description.

distribution-hash

Specifies the basis for the hash that the system uses as the frame distribution algorithm. The system uses the resulting hash to determine which interface to use for forwarding traffic.

When frames are transmitted on a trunk, they are distributed across the working member links. The distribution function ensures that the frames belonging to a particular conversation are neither mis-ordered nor duplicated at the receiving end. Distribution is done by calculating a hash value based on source and destination addresses carried in the frame and associating the hash value with a link. All frames with a particular hash value are transmitted on the same link, thereby maintaining frame order.

The options are:

dst-mac

Uses the destination MAC addresses to calculate the hash value.

`src-dst-mac`

Uses the source, destination, and MAC addresses to calculate the hash value.

`src-dst-ipport`

Uses the source and destination IP addresses and ports to calculate the hash value.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`id` Displays the ID of the trunk.

`interfaces`

Specifies the interfaces by name separated by spaces that you want to add to the trunk, delete from the trunk, or with which you want to replace all existing interfaces associated with the trunk.

`lACP` Specifies, when enabled, that the system supports the link aggregation control protocol (LACP), which monitors the trunk by exchanging control packets over the member links to determine the health of the links. If LACP detects a failure in a member link, it removes the link from the link aggregation. LACP is disabled by default, for backward compatibility.

`lACP-mode`

Specifies the operation mode for LACP if the lACP option is enabled for the trunk. The options are:

`active`

Specifies that the system periodically transmits LACP packets, regardless of the control value of the peer system.

`passive`

Specifies that the system periodically transmits LACP packets, unless the control value of the peer system is active.

`lACP-timeout`

Specifies the rate at which the system sends the LACP control packets. The default value is long.

The options are:

`long` Specifies that the system exchanges LACP packets every 30 seconds.

`short`

Specifies that the system exchanges LACP packets every second.

`link-select-policy`

Sets the LACP policy that the trunk uses to determine which member link (interface) can handle new traffic.

Link aggregation is allowed only when all the interfaces are operating at the same media speed and connected to the same partner aggregation system. When there is a mismatch among configured members due to configuration errors or topology changes (auto-negotiation), link selection policy determines which links become working members and form the aggregation.

The options are:

`auto` Specifies that the system chooses the lowest numbered operational link as the reference link. All the members that have the same media speed and are connected to the same partner as that of the reference link are declared as working members, and they are aggregated. The other configured members do not carry traffic.

`maximum-bandwidth`

Specifies that the system adds to the aggregation a subset of links that gives maximum aggregate bandwidth to the trunk.

`mac-address`

Specifies the media access control (MAC) address, which is associated with the trunk, in not case-sensitive hexadecimal colon notation, for example, 00:0b:09:88:00:9a.

`media`

Displays the media settings for the trunk.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

`stp` Enables or disables spanning tree protocols (STP). The default value is enabled.

If you disable STP, the system does not transmit or receive STP, RSTP, or MSTP packets on the trunk, and STP has no control over forwarding or learning on the trunk.

`stp-reset`

Resets STP, which forces a migration check.

`qinq-ethertype`

Specifies the ether-type value used for the packets handled on this trunk when it is a member in a QinQ vlan. The ether-type can be set to any string containing a valid hexadecimal 16 bits number, or any of the well known ether-types: 0x8100, 0x9100, 0x88a8. Default value is set to 0x8100.

working-mbr-count

Displays the number of working members associated with this trunk.

SEE ALSO

create, delete, edit, glob, list, modify, net interface, net vlan, net vlan-group, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2014-06-04 net trunk(1)

net tunnels endpoint

NAME

endpoint - Tunnel endpoint statistics

MODULE

net tunnels

SYNTAX

Show the statistics of tunnel remote endpoints within the net tunnel component using the syntax described in the following sections.

DISPLAY

show endpoint tunnel-name [name]

options:

remote-address [IP address]

MODIFY

reset-stats endpoint tunnel-name [name]

options:

remote-address [IP address]

DESCRIPTION

The command displays the statistics of tunnel remote endpoints. Currently, the command is only supported for network virtualization tunnels, e.g., VXLAN and NVGRE tunnels.

EXAMPLES

show endpoint tunnel-name t1

Display the statistics of remote endpoints of tunnel t1.

show endpoint tunnel-name t1 remote-address 10.10.0.2

Display the statistics of remote endpoint 10.10.0.2 of tunnel t1.

reset-stats endpoint tunnel-name t1

Reset the statistics of remote endpoints of tunnel t1.

reset-stats endpoint tunnel-name t1 remote-address 10.10.0.2

Reset the statistics of remote endpoint 10.10.0.2 of tunnel t1.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2016-09-23 net tunnels endpoint(1)

net tunnels etherip

NAME

etherip - Configures an EtherIP tunnel profile.

MODULE

net tunnels

SYNTAX

Configure the etherip component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

```
create etherip [name]
modify etherip [name]
options:
  app-service [[string] | none]
  defaults-from [name]
  description [string]
```

```
edit etherip [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list etherip
list etherip [ [ [name] | [glob] | [regex] ] ... ]
show running-config etherip
show running-config etherip [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete etherip [ all | [name] ]
```

DESCRIPTION

You can use the etherip component to create an EtherIP profile that you associate with a tunnel using the tunnel component. This will cause ethernet frames to be sent over the tunnel. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create etherip my_etherip
```

Creates an EtherIP profile called my_etherip.

```
list etherip all-properties
```

Displays all of the properties of all EtherIP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is etherip.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

net tunnels fec-stat

NAME

fec-stat - Displays FEC tunnels statistics.

MODULE

net tunnels

SYNTAX

Display statistics for the FEC tunnels using the syntax in the following section.

DISPLAY

show action

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the fec-stat component to display FEC tunnels statistics. The statistics details are described below:

name Specifies the FEC tunnel name.

profile

Specifies the FEC profile name used for the tunnel.

out_raw_packets

Specifies the number of FEC outgoing raw packets coming from the LAN.

out-raw-bits

Specifies the number of FEC outgoing raw bits coming from the LAN.

out_rdnd_packets

Specifies the number of FEC outgoing redundant packets sent to the WAN.

out_rdnd_bits

Specifies the number of FEC outgoing redundant bits sent to the WAN.

in_raw_packets

Specifies the number of incoming raw packets sent to the LAN.

in-raw-bits

Specifies the number of incoming raw bits sent to the LAN.

in_rdnd_packets

Specifies the number of FEC redundant packets incoming from the WAN.

in_rdnd_bits

Specifies the number of FEC redundant bits incoming from the WAN.

in_rdnd_lost

Specifies the number of FEC redundant packets lost incoming from the WAN.

in_raw_lost

Specifies the number of incoming from WAN raw packets lost.

rmt_in_rdnd_packets

Specifies the number of FEC redundant packets from a remote server when incoming from the WAN.

rmt_in_raw_packets

Specifies the number of raw packets from remote a server when incoming from the WAN.

rmt_in_rdnd_lost

Specifies the number of FEC redundant packets lost from a remote server when incoming from the WAN.

rmt_in_raw_lost

Specifies the number of raw packets lost from a remote server when incoming from the WAN.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, fec, net tunnels, net tunnels fec, net tunnels tunnel, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 net tunnels fec-stat(1)

net tunnels fec

NAME

fec - Configures a Forward Error Correction (FEC) profile.

MODULE

net tunnels

SYNTAX

Configure the fec component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create fec [name]

modify fec [name]

options:

app-service [[string] | none]

decode-idle-timeout [integer]

decode-max-packets [integer]

decode-queues [integer]

defaults-from [name]

description [string]

encode-max-delay [integer]

keepalive-interval [integer]

lzo [disabled | enabled]

repair-adaptive [disabled | enabled]

repair-packets [integer]

source-adaptive [disabled | enabled]

source-packets [integer]

udp-port [integer]

edit fec [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

=head2 DISPLAY

list fec

list fec [[name] | [glob] | [regex]] ...]

show running-config fec

show running-config fec [[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete fec [all | [name]]

DESCRIPTION

You can use the fec component to create a FEC profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create fec my_fec
```

Creates a FEC profile called my_fec.

```
list fec all-properties
```

Displays all of the properties of all of the FEC profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default

value is fec.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

decode-idle-timeout

Specifies the maximum waiting time for packets in decoding queues. Packets waiting longer than this time are discarded. Range is from 250 to 2000 milliseconds. The default value is 1500 milliseconds.

decode-max-packets

Specifies the maximum number of waiting packets in decoding queues. Range is from 200 to 8000. The default value is 512.

decode-queues

Specifies the number of decoding queues. Valid numbers are 8,16,32,64,128,256,512,1024. The default value is 32.

encode-max-delay

Specifies the maximum waiting time for packet aggregation. Range is from 500 to 5000 microseconds. The default value is 500 microseconds.

keepalive-interval

Specifies the interval between keepalive (statistical data) packets. Range is from 0 to 100 seconds. The default value is 5 seconds.

lzo Controls the use of the LZO compression algorithm to compress data packets. The default value is enabled.

repair-adaptive

Controls the use of the adaptive FEC repair technique to modify the number of redundant packets according to actual network conditions. The default value is enabled.

repair-packets

Specifies the number of redundant packets to add. Range is from 0 to 15. The default value is 15. This value should be less than or equal to the value specified for source-packets.

source-adaptive

Controls the use of the adaptive FEC source packets technique to reduce the number of packets for better MTU usage. The default value is enabled.

source-packets

Specifies the number of packets into which the system divides the aggregated payload. Range is from 1 to 15. The default value is 15.

udp-port

Specifies the local port for receiving FEC packets. The default value is 8288.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels ipip, net tunnels tunnel, net tunnels wccp, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2014-05-16 net tunnels fec(1)

net tunnels geneve

NAME

geneve - Configures a Geneve profile.

MODULE

net tunnels

SYNTAX

Configure the geneve component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

```
create geneve [name]
modify geneve [name]
options:
  app-service [[string] | none]
  defaults-from [ [name] | none]
  description [string]
  port [integer]
  flooding-type [none | multicast | multipoint]
```

```
edit geneve [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list geneve
list geneve [ [ [name] | [glob] | [regex] ] ... ]
show running-config geneve
show running-config geneve [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete geneve [ all | [name] ]
```

DESCRIPTION

You can use the geneve component to create a geneve profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create geneve my_geneve
```

Creates a geneve profile called my_geneve.

```
list geneve all-properties
```

Displays all the properties of all the geneve profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is geneve.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

port Specifies the local port for receiving geneve packets. The default is 6081.

flooding-type

Specifies the flooding type to use to transmit multicast, broadcast and unknown destination frames. The default is multipoint.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

net tunnels gre

NAME

gre - Configures a Generic Router Encapsulation (GRE) profile.

MODULE

net tunnels

SYNTAX

Configure the gre component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create gre [name]

modify gre [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

rx-csum [disabled | enabled]

tx-csum [disabled | enabled]

encapsulation [standard | nvgre]

flooding-type [none | multipoint]

edit gre [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list gre

list gre [[[name] | [glob] | [regex]] ...]

show running-config gre

show running-config gre [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete gre [all | [name]]

DESCRIPTION

You can use the gre component to create a GRE profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create gre my_gre
```

Creates a GRE profile called my_gre.

```
list gre all-properties
```

Displays all of the properties of all of the GRE profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is gre.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rx-csum

Specifies whether the system verifies the checksum on received packets. The default value is disabled.

tx-csum

Specifies whether the system includes a checksum on transmitted packets. The default value is disabled.

encapsulation

Specifies the flavor of GRE header to use for encapsulation. The default value is standard.

flooding-type

Specifies the flooding type to use to transmit broadcast and unknown destination frames. The default is none.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels ipip, net tunnels tunnel, net tunnels wccp, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2014-05-16 net tunnels gre(1)

net tunnels ipip

NAME

ipip - Configures an IP over IP (IPIP) profile.

MODULE

net tunnels

SYNTAX

Configure the ipip component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create ipip [name]

modify ipip [name]

options:

app-service [[string] | none]

defaults-from [name]

description [string]

proto [IPv4 | IPv6]

ds-lite [bool]

edit ipip [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ipip

list ipip [[[name] | [glob] | [regex]] ...]

show running-config ipip

show running-config ipip [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete ipip [all | [name]]

DESCRIPTION

You can use the ipip component to create an IPIP profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

create ipip my_ipip

Creates an IPIP profile called my_ipip.

list ipip all-properties

Displays all of the properties of all of the IPIP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is ipip.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the partition within which this component resides.

proto

Specifies the next header protocol. The default value is IPv4.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ds-lite

Specifies whether the profile is used for a DS-lite deployment. When enabled, an augmented flow lookup is made using the IPv6 address in the outer header in addition to the inner header addresses for packets coming over this tunnel. The default value is disabled.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels gre, net tunnels tunnel, net tunnels wccp, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2014-03-18 net tunnels ipip(1)

net tunnels ipsec

NAME

ipsec - Configures an IPsec profile.

MODULE

net tunnels

SYNTAX

Configure the ipsec component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create ipsec [name]

modify ipsec [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

traffic-selector [name]

edit ipsec [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

```
DISPLAY
list ipsec
list ipsec [ [ [name] | [glob] | [regex] ] ... ]
show running-config ipsec
show running-config ipsec [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

```
DELETE
delete ipsec [ all | [name] ]
```

DESCRIPTION

You can use the ipsec component to create an ipsec profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create ipsec my_ipsec
```

Creates an IPsec profile called my_ipsec.

```
list ipsec all-properties
```

Displays all the properties of all the IPsec profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is ipsec.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

traffic-selector

Specifies the IPsec traffic selector for the IPsec tunnel.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013. All rights reserved.

BIG-IP 2015-11-17 net tunnels ipsec(1)

net tunnels lw4o6

NAME

lw4o6 - Configures a LW4o6 tunnel profile.

MODULE

net tunnels

SYNTAX

Configure the lw4o6 component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

```
create lw4o6 [name]
modify lw4o6 [name]
options:
  app-service [[string] | none]
  defaults-from [ [name] | none]
  description [string]
  lwtbl-file [string]
  psid-length [integer]
  all-protocols-pass [enabled | disabled]

edit lw4o6 [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list lw4o6
list lw4o6 [ [name] | [glob] | [regex] ] ... ]
show running-config lw4o6
show running-config lw4o6 [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete lw4o6 [ all | [name] ]
```

DESCRIPTION

You can use the lw4o6 component to create a LW4o6 profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create lw4o6 my_lw4o6 lwtbl-file name
```

Creates a LW4o6 profile called my_lw4o6, lwtbl-file name have to be created before. For more information about creating a lwtbl-file see sys file lwtunneltbl.

```
list lw4o6 all-properties
```

Displays all the properties of all the LW4o6 profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is lw4o6.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

lwtbl-file

Specifies the existing lw4o6 configuration file.

psid-length

Specifies the length of Port Set ID (PSID). The default value is 0.

all-protocols-pass

When enabled and PSID is 0, permits passthrough for all IP sub-protocols. If PSID != 0, value have to be disabled. If it is disabled only TCP/UDP/ICMP traffic can pass lw4o6 tunnel. The default value is disabled."

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015. All rights reserved.

BIG-IP 2016-01-07 net tunnels lw4o6(1)

net tunnels map

NAME

map - Configures a MAP tunnel profile.

MODULE

net tunnels

SYNTAX

Configure the map component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create map [name]

modify map [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

ip6-prefix [ipv6 address/prefix length]

ip4-prefix [ipv4 address/prefix length]

ea-bits-length [integer]

port-offset [integer]

edit map [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list map

list map [[[name] | [glob] | [regex]] ...]

show running-config map

show running-config map [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete map [all | [name]]

DESCRIPTION

You can use the map component to create a MAP profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create map my_map
```

Creates a MAP profile called my_map.

```
list map all-properties
```

Displays all the properties of all the MAP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is map.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression

syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ip6-prefix

Specifies the IPv6 Prefix using CIDR notation, such as 2014::/48. The default prefix length is 48.

ip4-prefix

Specifies the IPv4 Prefix using CIDR notation, such as 192.0.0.0/8. The default prefix length is 8.

ea-bits-length

Specifies the length in bits of the EA (Embedded Address) of the MAP domain. The default is 32 (IPv4 prefix 24 bits + PSID 8 bits).

port-offset

Specifies the port offset bits length of the MAP domain. The default is 6.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015. All rights reserved.

BIG-IP 2016-03-14 net tunnels map(1)

net tunnels ppp

NAME

ppp - Configures a PPP profile.

MODULE

net tunnels

SYNTAX

Configure the ppp component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create ppp [name]

modify ppp [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

lcp-echo-failure [integer]

lcp-echo-interval [integer]

vj [disabled | enabled]

edit ppp [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list ppp

list ppp [[name] | [glob] | [regex]] ...]

show running-config ppp

show running-config ppp [[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete ppp [all | [name]]

DESCRIPTION

You can use the ppp component to create a ppp profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create ppp my_ppp
```

Creates a PPP profile called my_ppp.

```
list ppp all-properties
```

Displays all the properties of all the PPP profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is ppp.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

lcp-echo-failure

Specifies the number of consecutive PPP LCP echo messages that must go unanswered for the server to drop PPP connection. For example, if the server sends number of consecutive PPP LCP Echo Request messages that go unanswered (by Echo Reply), it will close the PPP connection. The default value is 4.

lcp-echo-interval

Specifies the interval, in seconds, between the PPP LCP Echo Request messages that the server sends to the peer (client). The default value is 30.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

vj Specifies whether the system uses Van Jacobson Header Compression (also known as VJ compression, or just Header Compression), which is an option in most versions of PPP. VJ is a data compression protocol described in RFC 1144, specifically designed by Van Jacobson to improve TCP/IP performance over slow serial links. The default value is disabled.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2014-05-16 net tunnels ppp(1)

net tunnels tcp-forward

NAME

tcp-forward - Configures a tcp-forward tunnel profile.

MODULE

net tunnels

SYNTAX

Configure the tcp-forward component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

```
create tcp-forward [name]
modify tcp-forward [name]
```

options:
app-service [[string] | none]
defaults-from [[name] | none]
description [string]

edit [[name] ...]
options:
all-properties
non-default-properties

DISPLAY
list tcp-forward
list tcp-forward [[[name] ...]
show running-config tcp-forward
show running-config tcp-forward [[name] ...]
options:
all-properties
app-service
non-default-properties
one-line
partition

DELETE
delete tcp-forward [all | [name]]

DESCRIPTION
You can use the tcp-forward component to create a TCP-FORWARD profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES
create tcp-forward my_tcp_forward

Creates a TCP_FORWARD profile called my_tcp_forward.

list tcp-forward all-properties

Displays all the properties of all the TCP-FORWARD profiles.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from
Specifies the existing profile from which the system imports settings for the new profile. The default value is tcp-forward.

description
User defined description.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition
Displays the administrative partition within which this component resides.

SEE ALSO
create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015, 2017. All rights reserved.

BIG-IP 2018-10-20 net tunnels tcp-forward(1)

net tunnels tunnel

NAME
tunnel - Configures a tunnel.

MODULE
net tunnels

SYNTAX
Configure the tunnel component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create tunnel [name]

modify tunnel [name]

options:

app-service [[string] | none]
auto-lasthop [default | enabled | disabled]
description [string]
local-address [ip address]
secondary-address [ip address]
mode [bidirectional | inbound | outbound]
mtu [integer]
use-pmtu [enabled | disabled]
profile [name]
remote-address [ip address]
traffic-group [[traffic group] | none]
tos [integer]
transparent [enabled | disabled]
idle-timeout [integer]
key [integer]

edit tunnel [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list tunnel

list tunnel [[[name] | [glob] | [regex]] ...]

show running-config tunnel

show running-config tunnel [[[name] | [glob] | [regex]] ...]

options:

all-properties
app-service
non-default-properties
one-line
partition

DELETE

delete tunnel [all | [name]]

DESCRIPTION

You can use the tunnel component to configure a tunnel.

EXAMPLES

```
create tunnel my_tunnel local-address 10.10.10.1 remote-address 11.11.11.1 profile gre
```

Creates a tunnel named my_tunnel between the local IP address 10.10.10.1 and the remote IP address 11.11.11.1.

```
list tunnel all-properties
```

Displays all of the properties of all of the tunnels.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot

modify or delete the object. Only the application service can modify or delete the object.

auto-lasthop

When enabled, specifies that the system returns packets to the MAC address from which they were sent. The default setting is default, which specifies that the system uses the default route to send back the request.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

if-index

Displays the index assigned to this tunnel. It is a unique identifier assigned for all objects displayed in the SNMP IF-MIB.

local-address

Specifies a local IP address. This option is required.

secondary-address

Specifies a secondary non-floating IP address when the local-address is set to a floating address. Currently this setting is supported by NVGRE tunnels only.

mode Specifies how the tunnel carries traffic. The default value is bidirectional.

mtu Specifies the maximum transmission unit (MTU) of the tunnel. The default value is 0. When the MTU is set to the default value (of 0), the MTU of the tunnel is computed by the system and is set to the MTU size of the underlying interface minus the encapsulation overhead introduced by the tunneling protocol. The valid range is 0 - 65535.

use-pmtu

Enables or disables the tunnel to use the PMTU (Path MTU) information provided by ICMP NeedFrag error messages. If enabled and the tunnel MTU is set to 0, the tunnel will use the PMTU information. If enabled and the tunnel MTU is fixed to a non-zero value, the tunnel will use the minimum of PMTU and MTU. If disabled, the tunnel will use fixed MTU, or calculate its MTU using tunnel encapsulation configurations.

name Specifies a unique name for the component. This option is required for the commands create, and modify.

partition

Displays the administrative partition within which the component resides.

profile

Specifies the profile that you want to associate with the tunnel. This option is required for the create command.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

remote-address

Specifies a remote IP address. This value is required for the commands create and modify.

traffic-group

Specifies a traffic-group for use with the tunnel. Traffic group determines the ConfigSync behavior of the tunnel object.

tos Specifies a value for insertion into the Type of Service (ToS) octet within the IP header of the encapsulating header of transmitted packets. The default value is preserve. The possible values are 0 (zero) - 255.

transparent

Enables or disables the tunnel to be transparent. If enabled, the user can inspect and/or manipulate the encapsulated traffic flowing through the BIG-IP. A transparent tunnel terminates a tunnel while presenting the illusion that the tunnel transits the device unperturbed i.e. the BIG-IP appears like an intermediate router that simply routes IP traffic through the device. The default value is disabled.

idle-timeout

Specifies an idle timeout for wildcard tunnels in seconds. This setting specifies the number of seconds that a wildcard tunnel connection is idle before the connection is eligible for deletion. The default value is 300 seconds.

key The key field may represent different values depending on the type of the tunnel. For VXLAN it represents the Virtual Network Identifier (VNI). The default value is 0.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels gre, net tunnels ipip, net tunnels wccp, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 net tunnels tunnel(1)

net tunnels v6rd

NAME

v6rd - Configures a 6RD profile.

MODULE

net tunnels

SYNTAX

Configure the v6rd component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create v6rd [name]

modify v6rd [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

v6rdprefix [IPv6 address]

v6rdprefixlen [integer]

ipv4prefix [IPv4 address]

ipv4prefixlen [integer]

edit v6rd [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list v6rd

list v6rd [[name] | [glob] | [regex]] ...]

show running-config v6rd

show running-config v6rd [[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete v6rd [all | [name]]

DESCRIPTION

You can use the v6rd component to create a v6rd profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

create v6rd my_v6rd

Creates a 6RD profile called my_v6rd.

list v6rd all-properties

Displays all the properties of all the 6RD profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is v6rd.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

v6rdprefix

Specifies the IPv6 prefix for 6rd domain.

v6rdprefixlen

Specifies the IPv6 prefix length of the 6rd domain. The default is 56.

ipv4prefix

As an extension not mentioned in the RFC5969, it specifies the IPv4 prefix for the Customer-Edge (CE) devices of a 6RD domain at a Border-Relay (BR) in case that the subnet prefixes used by the 6RD devices do not share the same IPv4 prefix. If they do, there is no need to configure this parameter. The default value is 0.0.0.0.

ipv4prefixlen

Also noted as IPv4MaskLen in RFC5969, it specifies the number of identical high-order bits shared by all CE and BR IPv4 addresses in a given 6RD domain. The valid range is from zero to 32. It is a required value for create. It defaults to zero, i.e. the full ipv4 address must be encapsulated.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

net tunnels vxlan

NAME

vxlan - Configures a VXLAN profile.

MODULE

net tunnels

SYNTAX

Configure the vxlan component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

create vxlan [name]

modify vxlan [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

port [integer]

flooding-type [none | multicast | multipoint | replicator]

encapsulation-type [vxlan | vxlan-gpe]

edit vxlan [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list vxlan

list vxlan [[[name] | [glob] | [regex]] ...]

show running-config vxlan

show running-config vxlan [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

DELETE

delete vxlan [all | [name]]

DESCRIPTION

You can use the vxlan component to create a vxlan profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

create vxlan my_vxlan

Creates a VXLAN profile called my_vxlan.

list vxlan all-properties

Displays all the properties of all the VXLAN profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is vxlan.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which this component resides.

port Specifies the local port for receiving VXLAN packets. The default is 4789.

flooding-type

Specifies the flooding type to use to transmit multicast, broadcast and unknown destination frames. The default is multicast.

encapsulation-type

Specifies whether the VXLAN header is formatted according to RFC 7348 (vxlan) or with the Generic Protocol Extension (vxlan-gpe). The default is vxlan.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 net tunnels vxlan(1)

net tunnels wccp

NAME

wccp - Configures a Web-cache coordination protocol (WCCP) GRE profile.

MODULE

net tunnels

SYNTAX

Configure the wccp component within the net tunnels module using the syntax in the following sections.

CREATE/MODIFY

```
create wccp [name]
modify wccp [name]
options:
  app-service [[string] | none]
  defaults-from [name]
  description [string]
  rx-csum [disabled | enabled]
  tx-csum [disabled | enabled]
  wccp-version [1 | 2]
```

```
edit wccp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list wccp
list wccp [ [ [name] | [glob] | [regex] ] ... ]
show running-config wccp
show running-config wccp [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  app-service
  non-default-properties
  one-line
  partition
```

DELETE

```
delete wccp [ all | [name] ]
```

DESCRIPTION

You can use the wccp component to create a WCCP GRE profile that you associate with a tunnel using the tunnel component. For more information about creating a tunnel see net tunnel.

EXAMPLES

```
create wccp my_wccp_gre
```

Creates a WCCP GRE profile called my_wccp_gre.

```
list wccp all-properties
```

Displays all of the properties of all of the WCCP GRE profiles.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the existing profile from which the system imports settings for the new profile. The default value is wccpgre.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

rx-csum

Specifies whether the system verifies the checksum on received packets. The default value is disabled.

tx-csum

Specifies whether the system includes a checksum on transmitted packets. The default value is disabled.

wccp-version

Specifies the version of WCCP that the system uses. The default value is 2.

SEE ALSO

create, delete, edit, glob, list, modify, net tunnels gre, net tunnels ipip, net tunnels tunnel, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2014-03-18 net tunnels wccp(1)

net vlan-allowed

NAME

vlan-allowed - Displays a list of available VLANs which can be used by the system.

MODULE

net

SYNTAX

Display the vlan-allowed component within the net module using the syntax shown in the following sections.

DISPLAY

show vlan-allowed

options:

field-fmt

DESCRIPTION

Displays a list of available VLANs which can be used by the system.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013. All rights reserved.

net vlan-group

NAME

vlan-group - Configures a VLAN group.

MODULE

net

SYNTAX

Modify the vlan-group component within the net module using the syntax shown in the following sections.

CREATE/MODIFY

```
create vlan-group [name]
modify vlan-group [name]
options:
  app-service [[string] | none]
  auto-lasthop [default | enabled | disabled ]
  bridge-in-standby [disabled | enabled]
  bridge-multicast [disabled | enabled]
  bridge-traffic [disabled | enabled]
  description [string]
  members
    [add | delete | replace-all-with] ] {
[vlan name] ...
  }
  members [default | none]
  migration-keepalive [disabled | enabled]
  mode [opaque | translucent | transparent | virtual-wire]
  proxy-excludes
    [add | delete | replace-all-with] ] {
[ip address] ...
  }
  proxy-excludes [default | none]

edit vlan-group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list vlan-group
list vlan-group [ [ [name] | [glob] | [regex] ] ... ]
show running-config vlan-group
show running-config vlan-group
[ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

```
show vlan-group
show vlan-group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DELETE

```
delete vlan-group [name]
```

DESCRIPTION

The vlan-group component defines a VLAN group, which is a grouping of two or more VLANs belonging to the same IP network for the purpose of allowing Layer 2 packet forwarding between those VLANs.

The VLANs between which the packets are to be passed must be on the same IP network, and they must be grouped using the vlan-group component. For example: modify vlan-group network11 members add { internal external }.

EXAMPLES

```
create vlan-group my_vlan-group members add { vlan1 vlan2 }
```

Creates a VLAN group named my_vlan-group that consists of VLANs named vlan1 and vlan2.

```
modify vlan-group proxy-excludes add { 10.10.10.1 }
```

Sets the global VLAN group proxy exclusion list.

delete vlan-group my_vlan-group

Deletes the VLAN group named my_vlan-group.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

bridge-traffic

When enabled, specifies that the VLAN group forwards all frames, including non-IP traffic. The default value is disabled.

bridge-in-standby

When enabled, specifies that the VLAN group forwards packets, even when the system is the standby unit in a redundant system. This option is designed for deployments in which the VLAN group exists on only one of the units. If that does not match your configuration, using this option may cause adverse effects. The default value is disabled.

bridge-multicast

When enabled, allows bridging of non-Internet Protocol (IP) Address Resolution Protocol (ARP) multicast frames across a VLAN group. An example of when you might want to use this option is when you are implementing the Spanning Tree Protocol (STP).

description

User-defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

if-index

Displays the index assigned to this VLAN group. It is a unique identifier assigned for all objects displayed in the SNMP IF-MIB.

members

The names of the VLANs that you want to add to or delete from the VLAN group.

migration-keepalive

Specifies whether the system will send keepalive frames (TCP keepalives and empty UDP packets depending on the connection type) when a node is moved from one VLAN group member to another VLAN group member for all existing connections that the system has to that node.

mode Specifies the level of exposure of remote MAC addresses within VLAN groups. The default value is translucent.

The options are:

virtual-wire

Use this option to create a Layer 2 bridge that only forwards traffic between two configured members. Traffic forwarded by such a VLAN group keeps intact the Ethernet header from ingress to egress, thus making this device transparent. A VLAN group configured to be virtual-wire is restricted to two member VLANs.

opaque

Use this option when you have a Cisco router in the network sending CDP packets to the system. Because opaque VLAN groups require a source and destination MAC address, and CDP packets do not contain a source and destination MAC address, the CDP packets are not forwarded through the VLAN group. This mode changes the MAC address to the MAC address assigned to the VLAN group, a proxy ARP with Layer 3 forwarding.

translucent

Uses the real MAC address of the requested host with the locally unique bit toggled. Specifies that the system uses Layer 2 forwarding with locally-unique bit, toggled in ARP response across VLANs.

transparent

Leaves the MAC address unchanged by the traffic management system. Specifies that the system uses Layer 2 forwarding with the original MAC address of the remote system preserved across VLANs.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

proxy-excludes

Specifies the IP addresses that you want to include in the proxy ARP exclusion list. If you use VLAN groups, you must configure a proxy ARP forwarding exclusion list. F5 Networks recommends that you configure this feature if you use VLAN groups with a redundant system. The reason is that both units need to communicate directly with their gateways and the back-end nodes. Creating a proxy ARP exclusion list prevents traffic from being proxied through the active unit due to proxy ARP. This traffic needs to be sent directly to the destination, not proxied.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, modify, net interface, net self, net vlan, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2017-07-06 net vlan-group(1)

net vlan

NAME

vlan - Configures a virtual local area network (VLAN).

MODULE

net

SYNTAX

Modify the vlan component within the net module using the syntax shown in the following sections.

CREATE/MODIFY

create vlan [name]

modify vlan [name]

options:

app-service [(string) | none]

auto-lasthop [default | enabled | disabled]

description [string]

failsafe [disabled | enabled]

failsafe-action [failover | failover-restart-tm | reboot | restart-all]

failsafe-timeout [integer]

fwd-mode [I3 | passive | virtual-wire | none]

interfaces

[add | delete | modify | replace-all-with] {

[name] ... {

[tagged | untagged]

tag-mode [customer | service | double | none]

}

}

interfaces none

learning [disable-drop | disable-forward | enable-forward]

mtu [integer]

sflow {

options:

poll-interval [integer]

poll-interval-global [no | yes]

sampling-rate [integer]

sampling-rate-global [no | yes]

}

source-checking [disabled | enabled]

tag [integer | 4096]

customer-tag [(string) | none]

cmp-hash [default | dst-ip | src-ip | ipport]

dag-tunnel [outer | inner]

dag-round-robin [disabled | enabled]

hardware-syncookie [disabled | enabled]

syn-cache-threshold [integer]

syn-flood-rate-limit [integer]

edit vlan [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list vlan

list vlan [[[name] | [glob] | [regex]] ...]

show running-config vlan

show running-config vlan

[[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

show vlan

show vlan [[[name] | [glob] | [regex]] ...]

options:

all-properties

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DELETE

delete vlan [name]

DESCRIPTION

VLANs are part of the configuration of the BIG-IP(r) network components. VLANs can be based on either ports or tags. When creating a VLAN, a tag value for the VLAN is automatically chosen unless you specify a tag value on the command line.

VLANs can have both tagged and untagged interfaces. You can add an interface to multiple VLANs as a tagged interface. You can add an interface to a single VLAN as an untagged interface. The tagged traffic can be single tagged and double tagged.

Note: To reset the statistics that display when you use the command sequence show vlan, you must reset the statistics for the trunks and interfaces associated with the VLAN.

EXAMPLES

```
create vlan my_vlan interfaces add { 1.2 1.3 1.4 }
```

Create the VLAN my_vlan that includes the interfaces 1.2, 1.3, and 1.4.

```
delete vlan my_vlan
```

Delete the VLAN named my_vlan.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User-defined description.

failsafe

Enables a fail-safe mechanism that causes the active cluster to fail over to a redundant cluster when loss of traffic is detected on a VLAN, and traffic is not restored during the failover timeout period for that VLAN. The default value is disabled.

When you set the VLAN failsafe option to enabled, the default failsafe-action value is restart-all. Therefore, when the fail-safe mechanism is triggered, all the daemons are restarted and the unit fails over.

failsafe-action

Specifies the action for the system to take when the fail-safe mechanism is triggered. The default value is failover-restart-tm.

failsafe-timeout

Specifies the number of seconds that an active unit can run without detecting network traffic on this VLAN before it starts a failover. The default value is 90 seconds.

fwd-mode

Displays the current forwarding mode which is derived from the vlan member port-fwd-mode property. This property is read-only and cannot be modified. See "net interface" for details on port-fwd-mode.

The options are:

I3 The VLAN consists of interface member(s) with port-fwd-mode set to I3.

passive

The VLAN consists of interface member(s) with port-fwd-mode set to passive.

virtual-wire

The VLAN consists of interface member(s) with port-fwd-mode set to virtual-wire.

none The VLAN has no interface member.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

if-index

Displays the index assigned to this VLAN. It is a unique identifier assigned for all objects displayed in the SNMP IF-MIB.

interfaces

Specifies a list of tagged or untagged interfaces and trunks that you want to configure for the VLAN. Use tagged interfaces or trunks when you want to assign a single interface or trunk to multiple VLANs.

A tagged interface is one that you assign to a VLAN in a way that causes the system to add a VLAN tag into the header of any frame passing through that interface or trunk.

A trunk is a combination of two or more interfaces and cables configured as one link.

tag-mode

Specifies the tag mode of the interface or trunk associated with. The default value is none.

The available values are:

customer

Specifies tag-mode setting for vlan members that are facing customer network and carry single tagged traffic.

service

Specifies tag-mode setting for vlan members that are facing the service provider networks and carry single tagged traffic.

double

Specifies tag-mode setting for vlan members that are facing the service provider networks and carry double tagged traffic.

none Specifies no tag-mode setting.

learning

Specifies whether switch ports placed in the VLAN are configured for switch learning, forwarding only, or dropped. The default value is enable-forward.

mtu Sets a specific maximum transition unit (MTU) for the VLAN. The default value is 1500. This value does not include the layer2 header.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

sflow

Specifies sFlow settings for the VLAN:

poll-interval

Specifies the maximum interval in seconds between two pollings. The default value is 0. To enable this setting, you must also set the poll-interval-global setting to no.

poll-interval-global

Specifies whether the global VLAN poll-interval setting, which is available under sys sflow global-settings module, overrides the object-level poll-interval setting. The default value is yes.

The available values are:

no Specifies to use the object-level poll-interval setting.

yes Specifies to use the global VLAN poll-interval setting.

sampling-rate

Specifies the ratio of packets observed to the samples generated. For example, a sampling rate of 2000 specifies that 1 sample will be randomly generated for every 2000 packets observed. The default value is 0. To enable this setting, you must also set the sampling-rate-global setting to no.

sampling-rate-global

Specifies whether the global VLAN sampling-rate setting, which is available under sys sflow global-settings module, overrides the object-level sampling-rate setting. The default value is yes.

The available values are:

no Specifies to use the object-level sampling-rate setting.

yes Specifies to use the global VLAN sampling-rate setting.

source-checking

Specifies that only connections that have a return route in the routing table are accepted. The default value is disabled.

tag Specifies a number that the system adds into the header of any frame passing through the VLAN. The value can be 1 through 4094, or 4096. The default is to not use this option, and the system assigns a tag number between 1 to 4094. A VLAN with the special tag 4096 is not used in the packet processing path; rather it assists with virtual-wire configuration, and such VLANs can only have interfaces with the port-fwd-mode property set to virtual-wire.

customer-tag

Specifies a number that the system adds into the header of any double tagged frame passing through the VLAN. The value can be any of the following: 1 through 4094, or none. The default is none.

cmp-hash

Specifies how the traffic on the VLAN will be disaggregated. The traffic disaggregation on the VLAN can be based on source ip, dest ip, or L4 ports. The default cmp hash uses L4 ports.

dag-tunnel

Specifies whether the ip tunnel traffic on the VLAN should be disaggregated based on the inner ip header or outer ip header. The default value is outer.

dag-round-robin

Specifies whether intended stateless traffic on the VLAN should be disaggregated in a round-robin order instead of using static hash. The stateless traffic include nonIP L2 traffic and user-specified UDP protocols. The sys db variable dag.roundrobin.redag allows HSBs to round robin stateless traffic to remote HSBs/blades.

hardware-syncookie

Enables hardware syncookie mode on a VLAN. When enabled, the hardware per-VLAN SYN cookie protection will be triggered when the certain traffic threshold is reached on supported platforms. The default value is disabled.

syncache-threshold

Specifies the number of outstanding SYN packets on the VLAN that will trigger the hardware per-VLAN SYN cookie protection. The default value is set to 6000 packets.

syn-flood-rate-limit

Specifies the max number of SYN flood packets per second received on the VLAN before the hardware per-VLAN SYN cookie protection is triggered. The default value is set at 1000 packets per second.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, net interface, net self, net vlan-group, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2017-12-14 net vlan(1)

net wccp

NAME

wccp - Configures Web Cache Communication Protocol (WCCP) services.

MODULE

net

SYNTAX

Configure the wccp component within the net module using the syntax in the following sections.

CREATE/MODIFY

create wccp [name]

modify wccp [name]

options:

app-service [[string] | none]

cache-timeout [integer]

description [string]

services [add | delete | replace-all-with] {
[object identifier] {

options:

alt-hash-fields [dest-ip | dest-port | src-ip | src-port | none]

app-service [[string] | none]

custom-mask[integer]

hash-fields [dest-ip | dest-port | src-ip | src-port | none]

password [string | none]

port-type [none | dest | source]

ports [integer]

priority [integer]

protocol [tcp | udp]

redirection-method [gre | I2]

return-method [gre | I2]

routers [add | delete | replace-all-with] {
[ip address ...]

}

traffic-assign [hash | mask]

tunnel-local-address [ip address]

tunnel-remote-addresses [add | delete | replace-all-with] {
[ip address ...]

}

weight [integer]

}

}

edit wccp [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY
list wccp
list wccp [[[name] | [glob] | [regex]] ...]
show running-config wccp
show running-config wccp [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line

DELETE
delete wccp [name]

DESCRIPTION

You can use the wccp component to create and modify WCCPv2 service groups. WCCPv2 is a content-routing protocol developed by Cisco Systems, Inc., which provides a mechanism to redirect traffic flows in real time. A WCCP service in this context is a set of redirection criteria and processing instructions that the BIG-IP(r) system applies to any traffic that a router in the service group redirects to the BIG-IP system.

EXAMPLES

```
list wccp service-wccp all-properties
```

Displays the services and their attributes in the service group named service-wccp.

```
modify server-wccp cache-timeout 40
```

Changes the cache-timeout setting to 40 for the service group named server-wccp.

```
modify server-wccp services modify { 77 {weight 60} }
```

Changes the weight setting to 60 for the service identified as 77 in the service group named server-wccp.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

cache-timeout

Specifies the frequency of control messages between the system and the router. The range is from 1 to 60 seconds. The default value is 10.

description

User-defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

services

Specifies the service group identifier, a number between 51 and 255 that matches a service ID configured on the router.

Adds, deletes, or replaces a set of services. You can configure the following options for a service:

alt-hash-fields

Specifies to the router which traffic attributes to use to determine which BIG-IP system it should forward traffic to for load balancing (traffic-assign mask) : destination IP address (dest-ip) or destination port (dest-port) or source IP address (src-ip), or source port (src-port).

app-service

Specifies the name of the application service to which the service belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the service. Only the application service can modify or delete the service.

custom-mask

Specifies the hexadecimal mask used when traffic-assign mask; it should have not more than 6 bits ON.

hash-fields

Specifies to the router which traffic attributes to use to determine which BIG-IP system it should forward traffic to for load balancing (traffic-assign hash): destination IP address (dest-ip), destination port (dest-port), source IP address (src-ip), and/or source port (src-port).

object identifier

Specifies the service group identifier, a number between 51 and 255 that matches a service ID

configured on the router.

password
Specifies the password for MD5 authentication or none.

port-type
Specifies whether the WCCP interception of traffic is based on the destination port (dest) or source port (source), or is not specified (none). The default value is none.

ports
Specifies one or more ports (up to 8) on which traffic is redirected.

priority
Specifies the precedence of the service group relative to the other service groups. The range is from 1 to 255. The default value is 100.

protocol
Specifies the protocol of the traffic to be redirected: TCP (tcp) or UDP (udp). The default value is tcp.

redirection
Specifies the method the router uses to redirect traffic: GRE gre or L2 I2. The default value is gre.

return
Specifies the method used to return passthrough traffic to the router: GRE (gre) or L2 (I2). The default value is gre.

routers
Specifies the IP addresses of the WCCP-enabled routers that redirect traffic.

traffic-assign
Specifies whether load balancing is achieved by a hash algorithm or a mask. If you specify hash, specify one or more attributes using the option hash-fields. If you specify mask, specify one attribute using the option alt-hash-fields

tunnel-local-address
Specifies an IP address on the BIG-IP system to which the WCCP-enabled routers should redirect traffic. Specify a self IP address of an external VLAN on the BIG-IP system.

tunnel-remote-addresses
Specifies the Router Identifier IP address of the router that redirects traffic.

weight
Specifies the relative importance of this traffic in a load-balancing environment. The range is from 1 to 100. The default value is 50.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2013. All rights reserved.

BIG-IP 2017-05-18 net wccp(1)

pem

pem forwarding-endpoint

NAME
forwarding-endpoint - Configures forwarding endpoints for the Policy Enforcement Manager (PEM).

MODULE
pem

SYNTAX
Modify the forwarding-endpoint component within the pem module using the syntax shown in the following sections.

CREATE/MODIFY
create forwarding-endpoint [name]
modify forwarding-endpoint [name]
options:
app-service [[string] | none]
description [[string] | none]

```

persistence {
  options:
type [destination-ip | disabled | hash | source-ip]
fallback [destination-ip | disabled | source-ip]
hash-settings {
  options:
  algorithm [carp ]
  length [integer]
  offset [integer]
  source [tcl-snippet | uri]
  tcl-value [string]
}
}
pool [name]
snat-pool [name]
source-port [change | preserve | preserve-strict]
translate-address [disabled | enabled]
translate-service [disabled | enabled]

```

```

edit forwarding-endpoint [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

```

DISPLAY
list forwarding-endpoint
list forwarding-endpoint [ [ [name] | [glob] | [regex] ] ... ]
show running-config forwarding-endpoint
show running-config forwarding-endpoint [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition

```

```

DELETE
delete forwarding-endpoint [name]

```

Note: All references to the forwarding-endpoint must be removed before it can be deleted.

DESCRIPTION
forwarding-endpoint is used to specify PEM policy forwarding action(s).

Note: A valid LTM pool with at least one member must be pre-configured before creating a forwarding-endpoint. Please refer to ltm pool for more info about configuring LTM pools.

EXAMPLES
create forwarding-endpoint my_endpoint { pool my_pool snatpool my_snatpool source-port preserved translate-address enabled translate-service enabled }

Creates a Policy Enforcement Manager forwarding endpoint named my_endpoint.

```
delete forwarding-endpoint my_endpoint
```

Deletes the forwarding-endpoint named my_endpoint.

```
list forwarding-endpoint my_endpoint
```

Displays the properties of the forwarding-endpoint named my_endpoint.

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, user cannot modify or delete the object. Only the application service can modify or delete the object.

description
Specifies a user-defined description.

persistence
Allows to set a specific persistence method for the pool member selection.

fallback
Specifies the fallback persistence method so that it applies when default persistence fail.

The options are:

destination-ip
Map the destination ip address to a specific pool member so that subsequent traffic sent to this address is directed to the same pool member.

source-ip
Map the source ip address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.

disabled
Specifies that this feature is disabled.

hash-settings
Specifies the settings for the hash persistence method.

algorithm
Specifies the algorithm to calculate the hash value. Currently only the carp algorithm is available.

length
Specifies the length of the source string used to calculate hash value. Default value of length is 1024.

offset
Specifies the offset in bytes from start of the source string to calculate the hash value. Default value of offset is 0.

source
Specifies the source for the string value which is used to calculate hash value.

tcl-value
Specifies the tcl script snippet so that when this script is executed its result used to calculate the hash value.

type Specifies the persistence method.

The options are:

destination-ip
Map the destination ip address to a specific pool member so that subsequent traffic sent to this address is directed to the same pool member.

hash Map the hash value to a specific pool member so that subsequent traffic with the same hash value is directed to the same pool member.

source-ip
Map the source ip address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.

disabled
Specifies that this feature is disabled.

pool Specifies the name of an LTM pool where the traffic is going to be directed to. Is used in the PEM policy rule forwarding actions. Note that the pool must be pre-configured before it can be referenced by a forwarding action.

snat-pool
Specifies the name of an existing LTM SNAT pool (snatpool) that is used to translate the client IP address to one of the configured IP addresses in that SNAT pool. The Self-IP addresses of the BIG-IP system must not be included in the SNAT pool. The default value is none.

source-port
Specifies whether the system preserves the source port of the connection. The default value is preserve.

The options are:

change
Specifies that the system changes the source port. This setting is useful for obfuscating internal network address.

preserve
Specifies that the system preserves the value configured for the source port, unless the source port from a particular snat is already in use, in which case the system uses a different port.

preserve-strict
Specifies that the system preserves the value configured for the source port. If the port is in use, the system does not process the connection. F5 Networks recommends restricting the use of this setting to cases that meet at least one of the following conditions:

The port is configured for UDP traffic.

The system is configured for nPath routing or is running in transparent mode (that is, there is no translation of any other Layer 3 or Layer 4 field).

There is a one-to-one relationship between virtual IP addresses and node addresses, or clustered multiprocessing (CMP) is disabled.

translate-address
Specifies, when enabled, that the system translates the original destination address of the virtual server. When disabled, specifies that the system uses the address without translation. The default value is disabled.

translate-service
Note that translate-service is really translate-port. It specifies, when enabled, that the system translates the original destination port. When disabled, it specifies that the system uses the original destination port without translation. The default value is disabled.

SEE ALSO

create, delete, edit, glob, list, modify, pem interception-endpoint, pem listener, pem policy, pem profile
diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber,
pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2016-01-07 pem forwarding-endpoint(1)

pem global-settings analytics

NAME

Analytics - Configures the global settings that pertain to Analytics reporting for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the analytics component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

```
modify analytics {  
  options:  
    mode [disabled | enabled]  
    subscriber-aware [disabled | enabled]  
    logging {  
      hsl {  
        endpoint-id [log-publisher]  
      }  
    }  
}
```

```
edit analytics [ [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties
non-default-properties

DISPLAY

```
list analytics  
list analytics [ [ [name] | [glob] | [regex] ] ... ]  
show running-config analytics  
show running-config analytics [ [ [name] | [glob] | [regex] ] ... ]  
options:  
  all-properties  
  non-default-properties  
  one-line  
  partition
```

DESCRIPTION

You can use the analytics component to configure global settings for analytics reporting.

EXAMPLES

```
modify analytics mode disabled subscriber-aware disabled logging hsl endpoint-id hsl_endpoint
```

Enables the analytics reporting for PEM and configures logging endpoint as hsl_endpoint.

```
list analytics
```

Displays the configuration for analytics settings.

OPTIONS

mode Specifies the mode for analytics reporting. It can take enable or disable as value.

subscriber-aware

Specifies the subscriber awareness for analytics reporting. It can take enable or disable as value.

logging

You can configure the following option for logging.

hsl You can configure the following options for hsl endpoint.

endpoint-id

Specifies the endpoint name.

SEE ALSO

create, delete, edit, glob, list, modify, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2014. All rights reserved.

BIG-IP 2014-09-05 pem global-settings analytics(1)

pem global-settings gx

NAME

gx - Configures the global settings that pertain to gx reporting for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the gx component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

modify gx

options:

gx-immediate-reporting [disabled | enabled]

gx-usage-record-merge [disabled | enabled]

edit gx [...]

options:

all-properties

non-default-properties

DISPLAY

list gx

list gx [...]

show running-config gx

options:

all-properties

current-module

non-default-properties

one-line

partition

DESCRIPTION

You can use the gx component to configure global settings for gx reporting.

EXAMPLES

modify gx gx-immediate-reporting disabled

Disables immediate reporting for gx.

list gx

Displays the configuration of gx reporting.

OPTIONS

gx-immediate-reporting

Displays the value of gx-immediate-reporting option.

gx-usage-record-merge

Displays the value of gx-usage-record-merge option.

SEE ALSO

create, delete, edit, glob, list, modify, pem reporting format-script, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2019. All rights reserved.

BIG-IP 2019-04-25 pem global-settings gx(1)

pem global-settings hsl-flow

NAME

hsl-flow - Configures the global settings that pertain to hsl flow reporting for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the hsl-flow component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

modify hsl-flow

options:

format-version

report-on-start [disabled | enabled]

report-on-interim [disabled | enabled]

edit hsl-flow [...]

options:

all-properties

non-default-properties

DISPLAY

list hsl-flow

list hsl-flow [...]

show running-config hsl-flow

options:

all-properties

non-default-properties

one-line

partition

DESCRIPTION

You can use the hsl-flow component to configure global settings for hsl flow reporting.

EXAMPLES

modify hsl-flow report-on-interim disable

Disables the generation of interim reports during the flow.

list hsl-flow

Displays the configuration for hsl-flow.

OPTIONS

format-version

DEPRECATED since version 15.1.0. Displays the value of formatting version used by hsl flow reporting.

report-on-start

Displays the value of report-on-start option.

report-on-interim

Displays the value of report-on-interim option.

SEE ALSO

create, delete, edit, glob, list, modify, pem reporting format-script, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013. All rights reserved.

BIG-IP 2019-05-17 pem global-settings hsl-flow(1)

pem global-settings hsl-report

NAME

hsl-report - Configures the global settings that pertain to hsl report for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the hsl-report component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

```
modify hsl-report
options:
  format-version
```

```
edit [ ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list
list [ ... ]
show running-config
options:
  all-properties
  current-module
  non-default-properties
  one-line
  partition
```

DESCRIPTION

You can use the hsl-report component to configure global settings for hsl report.

EXAMPLES

```
modify hsl-report format-version
```

Modifies the format version used for HSL reporting.

```
list hsl-report
```

Displays the configuration of hsl report.

OPTIONS

```
format-version
Displays the value of formatting version used by hsl reporting.
```

SEE ALSO

create, delete, edit, glob, list, modify, pem reporting format-script, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013. All rights reserved.

BIG-IP 2019-05-20 pem global-settings hsl-report(1)

pem global-settings insert-content

NAME

insert-content - Configures the global settings that pertain to insert content for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the insert-content component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

```
modify insert-content {
options:
  max-duration [value]
}
```

```
edit insert-content [ [ [name] | [glob] | [regex] ] ... ]
options:
```

all-properties
non-default-properties

DISPLAY
list insert-content
list insert-content [[[name] | [glob] | [regex]] ...]
show running-config insert-content
show running-config insert-content [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DESCRIPTION
You can use the insert-content component to configure global settings for content insert functionality.

EXAMPLES
modify insert-content max-duration 3600

Modifies the content insertion max duration when action is throttled to apply once.

list insert-content

Displays the configuration for insert-content settings.

OPTIONS
mode Specifies the max duration for applying insert content action when frequency Once is used.

SEE ALSO
create, delete, edit, glob, list, modify, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2014. All rights reserved.

BIG-IP 2014-12-26 pem global-settings insert-content(1)

pem global-settings policy

NAME
policy - Configures the global settings that pertain to policy for Policy Enforcement Manager (PEM).

MODULE
pem global-settings

SYNTAX
Modify the policy component within the pem global-settings module using the syntax shown in the following sections.

MODIFY
modify policy
options:
mandatory-action-list [rules [add | delete | modify | replace-all-with] {
options:
report-flow
}
single-rule-match-mode [disabled | enabled]

edit policy [...]
options:
all-properties
non-default-properties

DISPLAY
list policy
list policy [...]
show running-config policy
options:
all-properties
non-default-properties
one-line
partition

DESCRIPTION

You can use the policy component to configure global settings for PEM policy.

EXAMPLES

```
modify policy single-rule-match-mode disable
```

Disables the single rule match mode for PEM policy.

```
list policy
```

Displays the configuration for policy.

OPTIONS

```
single-rule-match-mode
```

Displays the value of single-rule-match-mode option.

```
mandatory-action-list
```

Displays the list mandatory actions applied.

SEE ALSO

create, delete, edit, glob, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2019. All rights reserved.

BIG-IP 2019-04-25 pem global-settings policy(1)

pem global-settings quota-mgmt

NAME

quota-mgmt - Configures the global settings that pertain to quota management over Gy for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the quota-mgmt component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

```
modify quota-mgmt
```

options:

```
default-rating-group [rating-group-name]
```

```
service-context-id [string]
```

```
edit quota-mgmt [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list quota-mgmt
```

```
list quota-mgmt [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config quota-mgmt
```

```
show running-config quota-mgmt [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

```
partition
```

DESCRIPTION

You can use the quota-mgmt component to configure global settings for quota management over Gy.

EXAMPLES

```
modify quota-mgmt default-rating-group rg_grp_1 service-context-id 32251@3gpp.org
```

Configures rg_grp_1 as default rating group and service-context-id as 32251@3gpp.org. rg_grp_1 should be defined before.

```
list quota-mgmt
```

Displays the configuration for quota-mgmt.

OPTIONS

default-rating-group
Specifies the default rating group for quota management over Gy.

service-context-id
Specifies the service-context-id to be used for CCR message over Gy.

SEE ALSO

create, delete, edit, glob, list, modify, pem quota-mgmt rating-group, pem listener, pem policy, pem profile
diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber,
pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013. All rights reserved.

BIG-IP 2013-10-23 pem global-settings quota-mgmt(1)

pem global-settings session-mgmt-attributes

NAME

session-mgmt-attributes - Configures the global settings that pertain to session management attributes for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the session-mgmt-attributes component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

modify session-mgmt-attributes

options:

clear-on-nas-reboot [enabled | disabled]
flow-term-on-sess-delete [enabled | disabled]
create-event
mode [enabled | disabled]
update-event
mode [enabled | disabled]
delete-event
mode [enabled | disabled]

edit session-mgmt-attributes

options:

clear-on-nas-reboot [enabled | disabled]
flow-term-on-sess-delete [enabled | disabled]
create-event
mode [enabled | disabled]
update-event
mode [enabled | disabled]
delete-event
mode [enabled | disabled]

DISPLAY

list session-mgmt-attributes

list session-mgmt-attributes [name]

show running-config session-mgmt-attributes

show running-config session-mgmt-attributes [name]

options:

clear-on-nas-reboot
flow-term-on-sess-delete
create-event
mode
update-event
mode
delete-event
mode

DESCRIPTION

You can use the session-mgmt-attributes component to configure global settings for session management attributes.

EXAMPLES

modify session-mgmt-attributes clear-on-nas-reboot enabled

Enable clear-on-nas-reboot so that when accounting-on or accounting-off Radius message is received for a particular NAS address, which indicates NAS reboot, will trigger clearing all PEM sessions associated with the NAS address.

modify session-mgmt-attributes flow-term-on-sess-delete enabled

Enable flow-term-on-sess-delete so that flow will be terminated shortly after session is marked for deletion.

modify session-mgmt-attributes create-event mode enabled

Enables event generation for subscriber session creation.

modify session-mgmt-attributes update-event mode disabled

Disables event generation for subscriber session update.

modify session-mgmt-attributes delete-event mode enabled

Enables event generation for subscriber session deletion.

list session-mgmt-attributes

Displays the configuration for session-mgmt-attributes.

OPTIONS

clear-on-nas-reboot

Specifies clear on nas reboot config for session-mgmt-attributes. The default value is disabled.

flow-term-on-sess-delete

Specifies flow termination on session delete config for session-mgmt-attributes. The default value is disabled.

create-event

You can configure the following options for create-event.

mode Specifies whether to generate event for subscriber session creation. The default value is disabled.

update-event

You can configure the following options for update-event.

mode Specifies whether to generate event for subscriber session update. The default value is disabled.

delete-event

You can configure the following options for delete-event.

mode Specifies whether to generate event for subscriber session deletion. The default value is disabled.

SEE ALSO

create, delete, edit, glob, list, modify, pem quota-mgmt, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2016. All rights reserved.

BIG-IP 2018-12-04 pem global-settings session-mgmt-attributes(1)

pem global-settings subscriber-activity-log

NAME

subscriber-activity-log - Configures the global settings that pertain to subscriber activity log messages for Policy Enforcement Manager (PEM).

MODULE

pem global-settings

SYNTAX

Modify the subscriber-activity-log component within the pem global-settings module using the syntax shown in the following sections.

MODIFY

modify subscriber-activity-log

options:

dynamic-subscriber-ids [add | delete | modify | replace-all-with] {
[id_name ...]

```
}
dynamic-subscriber-ids [none]
publisher [name]
static-subscriber-ids [add | delete | replace-all-with] {
  [id_name ...]
}
static-subscriber-ids [default | none]
subscriber-ip-addresses [add | delete | modify | replace-all-with] {
  [ip address/prefixlen ...]
}
subscriber-ip-addresses [none]
```

edit subscriber-activity-log [[[name] | [glob] | [regex]] ...]

options:
all-properties
non-default-properties

reset-stats subscriber-activity-log

DISPLAY

```
list subscriber-activity-log
list subscriber-activity-log [ [ [name] | [glob] | [regex] ] ... ]
show running-config subscriber-activity-log
show running-config subscriber-activity-log [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

show subscriber-activity-log

options:
all-properties
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
default
field-fmt

DESCRIPTION

You can use the subscriber-activity-log component to monitor behavior of the subscribers in the troubleshooting mode by sending activity log messages to one or more destinations. You can add static and dynamic subscribers by IDs, or by subscriber IP addresses. The activity log messages contain the internal information exposing the subscribers behavior.

EXAMPLES

```
modify subscriber-activity-log publisher pub1 dynamic-subscriber-ids add { 4081112222 }
```

Adds dynamic subscriber 4081112222 to troubleshooting mode by sending activity log messages to all destinations defined in pub1.

```
list subscriber-activity-log
```

Displays the list of the subscribers in troubleshooting mode.

```
show subscriber-activity-log
```

Displays the logging statistics of the subscribers in troubleshooting mode.

```
reset-stats subscriber-activity-log
```

Resets the logging statistics of the subscribers in troubleshooting mode.

OPTIONS

dynamic-subscriber-ids
Specifies a list of dynamic subscriber IDs to be in troubleshooting mode.

publisher
Specifies the external logging publisher used to send activity log messages to one or more destinations.

static-subscriber-ids
Specifies a list of static subscriber IDs to be in troubleshooting mode.

subscriber-ip-addresses
Specifies a list of subscriber IP addresses to be in troubleshooting mode.

SEE ALSO

create, delete, edit, glob, list, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 pem global-settings subscriber-activity-log(1)

pem interception-endpoint

NAME

interception-endpoint - Configures interception endpoints for the Policy Enforcement Manager (PEM).

MODULE

pem

SYNTAX

Modify the interception-endpoint component within the pem module using the syntax shown in the following sections.

CREATE/MODIFY

create interception-endpoint [name]

modify interception-endpoint [name]

options:

app-service [[string] | none]

persistence [destination-ip | disabled | source-ip]

pool [name]

edit interception-endpoint [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list interception-endpoint

list interception-endpoint [[[name] | [glob] | [regex]] ...]

show running-config interception-endpoint

show running-config interception-endpoint [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete interception-endpoint [name]

Note: You must remove all references to an interception-endpoint before you can delete the interception-endpoint.

DESCRIPTION

You can use the interception-endpoint component to configure interception-endpoint definitions for the Policy Enforcement Manager. The interception-endpoint is used to clone all traffic. Note: Before you create a cloning-endpoint you have to create a valid pool. Please refer to ltm pool for more information about how to create a pool.

EXAMPLES

```
create interception-endpoint my_endpoint { pool pool1 }
```

Creates a Policy Enforcement Manager interception-endpoint named my_endpoint.

```
delete interception-endpoint my_endpoint
```

Deletes the interception-endpoint named my_endpoint.

```
list interception-endpoint my_endpoint
```

Displays the properties of the interception-endpoint named my_endpoint.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

persistence

Specifies the persistence that is based on either the source or destination IP addresses only.

pool Specifies the pool. It is mandatory to specify a pool when creating any interception-endpoint. Before you create an interception-endpoint you have to create a valid pool.

SEE ALSO

create, delete, edit, glob, list, modify, pem forwarding-endpoint, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013. All rights reserved.

BIG-IP 2013-03-21 pem interception-endpoint(1)

pem irule

NAME

irule - Configures an PEM iRule for traffic management system configuration.

MODULE

pem

SYNTAX

Configure the irule component within the pem module using the syntax shown in the following sections.

CREATE/MODIFY

```
create irule [name]
edit irule [name]
modify irule [ [name] | [glob] | [regex] ] ... ]
```

Note: When using tmsh, you can only create pem iRule using the editor, which starts when you use the create or edit commands. You cannot create an pem iRule directly on the command line. The vim editor applies the autoindent and smartindent options. You can toggle on/off paste mode using the F12 key.

Note: You can also edit user metadata associated with a pem iRule. See the example section for more information.

DISPLAY

```
list irule
list irule [ [name] | [glob] | [regex] ] ... ]
show running-config irule
show running-config irule [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

show irule
show irule [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DELETE

```
delete irule [name]
```

DESCRIPTION

You can apply pem iRules as an action when the traffic matches the filter criteria defined in pem policy rule. The syntax that you use to write pem iRules is based on the Tools Command Language (Tcl) programming standard. Thus, you can use many of the standard Tcl commands, plus a robust set of extensions that the BIG-IP(r) policy enforcement management system provides to help you customize the actions you want to apply to the traffic.

You cannot edit the system iRules that come with the BIG-IP system. However, you can open a system iRule in the editor and use it as a template to create a new rule.

To create a new pem iRule using a system rule as a template:

1. Enter the command sequence edit irule [system rule name]. tmsh opens the system rule in an editor.
 2. Change the name of the rule in the editor.
 3. Edit the rule and exit the editor.
- tmsh checks for syntax errors, and if there are none, it saves the new rule.

For more information about iRules(r), see <http://devcentral.f5.com/>.

EXAMPLES

```
list irule
```

Displays all iRules.

```
delete irule my_irule
```

Deletes the pem iRule named my_irule.

```
create irule my_irule {
```

```
priority 1
when PEM_POLICY {
}
}
```

Creates a pem iRule named my_irule with priority 1.

```
modify rule my_irule {
  when RULE_INIT {}
  metadata replace-all-with {
my_meta {
  persist false
  value "hello"
}
my_meta2 {
  persist false
  value "hello 2"
}
}
}
```

Modifies an existing pem iRule named my_irule by adding a new metadata and modifying an existing metadata.

The metadata attribute is the user defined key/value pair. Metadata has the following format:

```
metadata
[add | delete | modify] {
[metadata_name] {
  value [ "value content" ]
  persist [ true | false ]
}
}

modify irule my_irule {
  when RULE_INIT {}
  metadata delete { my_meta }
}
```

Deletes a metadata from a pem iRule.

OPTIONS

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the create, delete, and modify commands.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

metadata
Specifies the user-defined key/value pair associated with the rule. See the example section for usage format.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 pem irule(1)

pem listener

NAME

listener - Configures listeners for the Policy Enforcement Manager (PEM).

MODULE

pem

SYNTAX

Modify the listener component within the pem module using the syntax shown in the following sections.

```
CREATE/MODIFY
create listener [name]
modify listener [name]
options:
  app-service [[string] | none]
  description [string]
  profile-spm [name]
  profile-subscriber-mgmt [name]
  virtual-servers [name] [add | delete | replace-all-with] {
    [virtual_server_name ... ]
  }
}
```

```
edit listener [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
DISPLAY
list listener
list listener [ [ [name] | [glob] | [regex] ] ... ]
show running-config listener
show running-config listener [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
DELETE
delete listener [name]
```

Note: You must remove all references to a listener before you can delete the listener.

DESCRIPTION

You can use the listener component to configure listener definitions for the Policy Enforcement Manager.

EXAMPLES

```
create listener lis1 { profile-spm spm1 virtual-servers add {vs_tcp vs_udp} }
```

Creates a Policy Enforcement Manager listener named lis1.

```
delete listener lis1
```

Deletes the listener named lis1.

```
list listener lis1
```

Displays the properties of the listener named lis1.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

profile-spm

Specifies the spm profile name.

profile-subscriber-mgmt

Specifies the subscriber-mgmt profile name.

virtual-servers

Adds, deletes, or replaces a set of virtual servers, by specifying a virtual server name.

SEE ALSO

create, delete, edit, glob, list, modify, pem forwarding-endpoint, pem interception-endpoint, pem policy, pem profile diameter-endpoint, pem profile spm, pem profile subscriber-mgmt, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012. All rights reserved.

BIG-IP 2016-08-17 pem listener(1)

pem policy

NAME

policy - Configures policies for the Policy Enforcement Manager (PEM).

MODULE

pem

SYNTAX

Modify the policy component within the pem module using the syntax shown in the following sections.

CREATE/MODIFY

```
create policy [name]
modify policy [name]
options:
  description [string]
  status [enabled | disabled]
  transactional [enabled | disabled]
  rules [add | delete | modify | replace-all-with] {
    [rule_name ... ] {
      options:
    }
  }
  app-service [[string] | none]
  classification-filters [add | delete | modify | replace-all-with] {
    [filter_name ...] {
      options:
        app-service [[string] | none]
        application [application_name]
        category [category_name]
        operation [match | nomatch]
    }
  }
  dscp-marking-downlink [integer]
  dscp-marking-uplink [integer]
  dtos-tethering {
    options:
      dtos-detect [enabled | disabled]
      tethering-detect [enabled | disabled]
      report {
        dest {
        }
      }
      hsl {
        options:
          format-script [ [format_script_name] | none]
          publisher [ [publisher_name] | none ]
        }
      }
  }
  ran-congestion {
    options:
      detect [enabled | disabled]
      lowerthreshold-bw [integer]
      report {
        dest {
        }
      }
      hsl {
        options:
          format-script [ [format_script_name] | none]
          publisher [ [publisher_name] | none ]
        }
      }
  }
  flow-info-filters [add | delete | modify | replace-all-with] {
    [filter-name ...] {
      options:
        app-service [[string] | none]
        dscp-code [integer]
        dst-ip-addr [ip address/prefixlen]
        dst-port [port]
        from-vlan [vlan_name]
        l2-endpoint [disabled | vlan]
        operation [match | nomatch]
        ip-addr-type [IPv4 | IPv6 | any]
        proto [ tcp | udp | any]
        src-ip-addr [ip address/prefixlen]
        src-port [port]
    }
  }
  flow-info-filters [none]
  forwarding {
    options:
      endpoint [forwarding_endpoint_name]
      fallback-action [drop | continue]
      internal-virtual [name]
      icap-type [request | response | both | none]
      type [icap | pool | route-to-network | none]
  }
}
```

```

gate-status [enabled | disabled]
http-redirect {
  options:
    redirect-url [string]
    fallback-action [drop | continue]
}
intercept [intercept_endpoint_name]
l2-marking-downlink [integer]
l2-marking-uplink [integer]
tcp-optimization-downlink [string]
tcp-optimization-uplink [string]
tcp-analytics-enable [enabled | disabled]
modify-http-hdr {
  options:
    name [header_name]
    operation [insert | none | remove]
    value-content [header_value]
    value-type [string | tcl-snippet]
}
insert-content {
  options:
    duration [integer]
    frequency [always | once | once-every]
    position [append | prepend]
    tag_name [name]
    value-content [string]
    value-type [string | tcl-snippet]
}

precedence [integer]
deprecated since 15.0.0:
qoe-reporting {
  options:
    dest {
      hsl {
        options:
          format-script [ [format_script_name] | none]
          publisher [ [publisher_name] | none ]
        }
      }
}
reporting {
  options:
    dest {
      gx {
        options:
          application-reporting [enabled | disabled]
          monitoring-key [name]
        }
      hsl {
        options:
          publisher [name]
          format-script [name]
          session-reporting-fields
            [add | delete | replace-all-with] {
[reporting field ... ]
          }
          flow-reporting-fields
            [add | delete | replace-all-with] {
[reporting field ... ]
          }
          transaction-reporting-fields
            [add | delete | replace-all-with] {
[reporting field ... ]
          }
        }
      radius-accounting {
        options:
          radius-aaa-virtual [name]
        }
      sd {
        options:
          application-reporting [enabled | disabled]
          monitoring-key [name]
        }
      granularity [flow | session | transaction]
      interval [integer]
      transaction {
        http {
        options:
          hostname-len [integer]
          uri-len [integer]
          user-agent-len [integer]
        }
      }
    }
  volume {

```

```

    options:
downlink
total
uplink
}
}
quota {
    options:
        rating-group [name]
        reporting-level [rating-group | service-id]
}
qos-rate-pir-downlink [bwc policy name | none]-> [category name | none]
qos-rate-pir-uplink [bwc policy name | none]-> [category name | none]
service-chain [service chain endpoint name]
sfc-action {
    options:
        path-name [string]
        metadata-template [string]
}
tcl-filter [tcl-script]
url-categorization-filters [add | delete | modify | replace-all-with] {
    [filter_name ...] {
        options:
            category [category_name]
            operation [match | nomatch]
    }
}
}
}
rules [none]

```

```
edit policy [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
    all-properties
    non-default-properties
```

DISPLAY

```
list policy
list policy [ [ [name] | [glob] | [regex] ] ... ]
show running-config policy
show running-config policy [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
    all-properties
    non-default-properties
    one-line
    partition
```

```
show policy
show policy [name]
options:
    all-properties
    (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
    detail
    field-fmt
```

DELETE

```
delete policy [name]
```

Note: You must remove all references to a policy before you can delete the policy.

DESCRIPTION

You can use this policy component to configure the policy definitions on the Policy Enforcement Manager. A policy is a set of rules which are used to match traffic flow and apply actions. A rule has configuration for filters and actions. All configured filters must match before the actions can be applied to the traffic flow. There are four filters: classification-filter, url-category-filter, flow-info-filter, and tcl-filter. Classification-filter allows for matching the traffic based on the flow L7 features, such as a specific application (for example, Google Mail) or application category (for example, Web). URL-category-filter allows for matching the type of URL, such as adult content. Flow-info-filter allows for matching the traffic using L2-L4 flow parameters. Tcl-filter provides a customized method to match traffic flows using iRule commands. The actions can be steering or/and reporting. Steering allows the user to manipulate the traffic when all configured filters match the flow. The steering options can be forwarded (option forwarding), drop/pass(option gate-status), redirect(option http-redirect), or intercept(option intercept). Reporting allows the user to report the usage to different endpoints by different output formats. The reporting options can be gx or hsl. Policy attribute transactional allow policy enforcement for HTTP traffic for each transaction. Quota allows users to do quota management over Gy by specifying the rating group, which has all the parameters associated.

EXAMPLES

```
create policy my_policy rules add {
    rule_1 {
        flow-info-filters {
            flow_1 {
                dscp-code 8
            }
            flow_2 {
                dst-port 80
            }
        }
        forwarding {
```

```

    endpoint server1
    fallback-action continue
  }
}
precedence 1
}
rule_2 {
  reporting {
  dest {
    hsl {
    endpoint-id pem_hsl
    format-script fm1
    }
  }
  granularity flow
  volume {
    total 5000
  }
}
precedence 2
}
}
}

```

Creates a Policy Enforcement Manager policy named my_policy with two rules, rule_1 and rule_2. rule_1 defines the flow-info-filters so that when the flow with DSCP is 8 or destination port is 80, the traffic will be forwarded to server1. rule_2 defines a flow-based reporting rule which will send flow usage record to pem_hsl endpoint using format script defined in fm1 whenever total increases by 5000 bytes.

delete policy my_policy

Deletes the policy named my_policy.

list policy my_policy

Displays properties of the policy named my_policy.

OPTIONS

app-service

Specifies the name of the application service to which the policy belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the policy. Only the application service can modify or delete the policy.

description

User defined description.

transactional

Indicate the policy enable or disable policy enforcement for each HTTP transaction.

partition

Displays the administrative partition within which the policy resides.

rules

Adds, deletes, or replaces a set of rules, by specifying a rule name. If a rule by the specified name does not exist, it will be created. You can configure the following options for a rule:

app-service

Specifies the name of the application service to which the rule belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the rule. Only the application service can modify or delete the rule.

classification-filters

Adds, deletes, or replaces a set of classification-filters. You can configure the following options for a classification-filter.

app-service

Specifies the name of the application service to which the classification-filter belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the rule. Only the application service can modify or delete the classification-filter.

application

Specifies the name of the application where the rule applies to the traffic. The default value is none.

category

Specifies the name of the category of applications where the rule applies to the traffic. The default value is none.

operation

The options match and nomatch indicate the traffic flow must match or not match the condition specified in the classification filter. The default value is match.

dscp-marking-downlink

Specifies the action to modify the DSCP code in the downlink packet when the traffic flow matches the rule matching criteria. The range is 0 to 63, or pass-through. The default value is pass-through, indicating the DSCP code of the downlink packet will not be changed when the traffic flow matches the rule.

dscp-marking-uplink
Specifies the action to modify the DSCP code in the uplink packet when the traffic flow matches the rule matching criteria. The range is 0 to 63, or pass-through. The default value is pass-through, indicating the DSCP code of the uplink packet will not be changed when the traffic flow matches the rule.

dtos-tethering
Defines the device type & OS and tethering detection action and its options.

dtos-detect
Specifies the device type & OS detection to be enabled or disabled. Default is disabled

tethering-detect
Specifies the tethering detection to be enabled or disabled. Default is disabled

report
You can configure the following options for dtos and tethering reporting.

dest You can configure the following options for destination.

hsl You can configure the following options for hsl publisher.

publisher
Specifies the publisher name.

format-script
Specifies the format script name to format the HSL output string format.

ran-congestion
Detect congestion in the Radio Access Network.

detect
Enable or disable the ran congestion detection. Default is disabled.

lowerthreshold-bw
Configured lowerthreshold bandwidth for a session in kbps. Session bandwidth below this value will be marked as congested. Default is 1000kbps.

report
You can configure the following options for ran congestion reporting.

dest You can configure the following options for destination.

hsl You can configure the following options for hsl publisher.

publisher
Specifies the publisher name.

format-script
Specifies the format script name to format the HSL output string format.

flow-info-filters
Adds, deletes, or replaces a set of the flow-info-filters. The flow info filter defines the flow conditions (Layer 4) that the traffic should meet (or not meet) for this enforcement policy rule to apply. You can configure the following options for a flow-info-filter.

app-service
Specifies the name of the application service to which the flow-info-filter belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the rule. Only the application service can modify or delete the flow-info-filter.

dscp-code
Specifies the value of DSCP code which matches incoming traffic based on a value in the DSCP field in the IP header. The range is 0 to 63, or disabled. The default value is disabled, indicating that the DSCP code will not be used to filter the packet in the flow-info-filter.

dst-ip-addr
Specifies the destination IP address and prefix length that the rule applies to. The format is [ip address/prefixlen]. The default value is 0.0.0.0/0.

dst-port
Specifies the destination port against which the packet will be compared. The default value is any.

from-vlan
Specifies the name of the source vlan to match the ingress flow arriving from that vlan.

l2-endpoint
Specifies an L2 endpoint type to be used when matching the traffic flows. The default value is disabled, indicating that L2 endpoint is not used for matching the flows. You can configure the following options:

disabled
Flows are not matched based on the L2 endpoint specification.

vlan The vlan name specified in from-vlan is used to match the traffic flows.

operation

Specifies whether the rule applies to traffic that matches (match) or does not match (nomatch) the traffic flow defined here. The options are match and nomatch. The default value is match.

proto

Specifies the protocol that this rule applies to. The options are any, tcp, and udp. The default value is any.

ip-add-type

Specifies the ip address type (IPv4 or IPv6) that this rule applies to. The options are any, IPv4, and IPv6. The default value is any.

src-ip-addr

Species the source IP address and prefix length that the rule applies to. The format is [ip address/prefixlen]. The default value is 0.0.0.0/0.

src-port

Specifies the source port of the network you want the rule to affect. The default value is any.

forwarding

Manages the forwarding action and its attributes.

endpoint

Specifies the forwarding endpoint. The endpoint can be icap, pool or route-to-network. Depending on the type chosen flow can be steered to icap server, pool or to the network.

fallback-action

Specifies whether the connection should continue unchanged or should be dropped in the event the forwarding action fails for any reason. The options are: drop or continue, and the default is drop.

internal-virtual

Specifies the internal virtual server name if the type selected is icap.

icap-type

Defines the ICAP adaptation type: request only adaptation, request and response adaptation or both types of adaptations combined.

type Specifies the type of forwarding action.

gate-status

Specifies, when set to enabled, that the traffic can pass through the system without being changed. Set disabled to drop traffic that this rule applies to. The options are disabled and enabled. The default is enabled.

http-redirect

Manages the HTTP redirect action and its attributes.

redirect-url

Specifies the HTTP redirection URL.

fallback-action

Specifies whether the connection should continue unchanged or should be dropped in the event the forwarding action fails for any reason. The options are: drop or continue, and the default is drop.

intercept

Specifies the name of the intercept endpoint.

l2-marking-downlink

Set Layer-2 Quality of Service Marking in downlink traffic that matches a rule. Setting a L2 QoS Marking affects the packet delivery priority. The range is 0 to 7, or pass-through. The default value is pass-through, indicating the L2 QoS Marking of the packet will not be changed when the packet matches the rule.

l2-marking-uplink

Set Layer-2 Quality of Service Marking in uplink traffic that matches a rule. Setting a L2 QoS marking affects the packet delivery priority. The range is 0 to 7, or pass-through. The default value is pass-through, indicating the L2 QoS Marking of the packet will not be changed when the packet matches the rule.

tcp-optimization-uplink

Set tcp optimization profile to be applied to the uplink traffic that matches a rule. The profile name should be one from the common tcp profile list.

tcp-optimization-downlink

Set tcp optimization profile to be applied to the downlink traffic that matches a rule. The profile name should be one from the common tcp profile list.

tcp-analytics-enable

Specifies the action to enable tcp analytics when the traffic flow matches the rule matching criteria. The options are disabled and enabled. The default is disabled.

modify-http-hdr

Specifies the action to modify the HTTP header when the traffic flow matches the rule matching

criteria. You can configure the following options for modifying the HTTP header.

name Specifies the HTTP header name used by the operation option to modify the HTTP header.

operation

Specifies the operation used to modify the HTTP header. The options are insert, none, and remove. The default value is none which indicates that no HTTP header modifications will be made.

value-content

Specifies the HTTP header value content used by the operation option to modify the HTTP header. Based on the selected value-type option, the content format will be interpreted either as a string or a tcl snippet. Note: This field is applicable only when the operation option is set to insert.

value-type

Specifies the type of content format used in the value-content field. The options are string and tcl-snippet. The default value is string which indicates that the value-content field will be interpreted as a string.

insert-content

Specifies the action to insert content into the webpage.

duration

Specifies the periodicity of the insert action. Note: This value is useful only when the frequency is set to once-every.

frequency

Specifies the frequency of the insert content action. It can take values once, once-every, always.

The options are:

always

Specifies if the action need to be applied always on the matched flow.

once Specifies if the action need to be applied once per subscriber.

once-every

Specifies if the action need to be applied once-every time interval configured in duration per subscriber.

position

Specifies the position with respect to the tag name configured. It can take values append, prepend.

value-content

Specifies the value content to be inserted into the webpage. Based on the selected value-type option, the content format will be interpreted either as a string or a tcl-snippet.

value-type

Specifies the type of content format used in the value-content field. The options are string and tcl-snippet. The default value is string which indicates that the value-content field will be interpreted as a string.

tag_name

Specifies the tag name to which the content is either appended or prepended.

precedence

Specifies the precedence for the rule in relation to the other rules. The range is 1 to 4294967295 where 1 has the highest precedence. A rule with higher precedence is evaluated at a high priority. It is mandatory to specify precedence when creating a rule in a policy.

qoe-reporting

Deprecated since 15.0.0. You can configure the following options for Quality-of-Experience (QoE) reporting.

dest You can configure the following options for destination.

hsl You can configure the following options for hsl publisher.

publisher

Specifies the publisher name.

format-script

Specifies the format script name to format the HSL output string format.

reporting

You can configure the following options for reporting.

dest You can configure the following options for destination.

gx You can configure the following options for gx endpoint.

application-reporting

Specifies whether the application reporting is enabled. When it is enabled, the APPLICATION_START and APPLICATION_STOP Event-Triggers will be reported when the

application start/stop is detected. The default value is disabled.

monitoring-key

Specifies the monitoring-key.

hsl You can configure the following options for hsl endpoint.

publisher

Specifies the publisher.

format-script

Specifies the format script name to format the HSL output string format.

session-reporting-fields

Specifies the session fields and their order based on which messages should be published.

3gpp-parameters

Reports the 3gpp-parameters of the session subscriber.

application-id

Reports the application/category ID that is classified for this session.

called-station-id

Reports the called station ID of the session subscriber.

calling-station-id

Reports the calling station ID of the session subscriber.

concurrent-flows

Reports the number of concurrent flows of this session.

downlink-volume

Reports the aggregate incoming bytes for the traffic associated with this session.

duration-seconds

Reports the total duration of all the flows belonging to the traffic associated with this session.

last-record-sent

Reports the time (seconds) when sending the last record.

new-flows

Reports the number of new flows associated with this session since last record.

observation-time-seconds

Reports the timestamp of the record.

record-reason

Reports the reason for sending the record.

record-type

Reports the reporting record type as 3 : session based record.

report-id

Reports the reporting module ID.

report-version

Reports the format version of this record.

subscriber-id

Reports the subscriber ID that of this session.

subscriber-id-type

Reports the ID type of the subscriber of this session.

successful-transactions

Reports the total number of successful transactions associated with this session.

terminated-flows

Reports the total number of terminated flows during this session.

timestamp-msec

Reports the time stamp on this record in milli-seconds.

total-transactions

Reports the total number of transactions of this session.

uplink-volume

Reports the aggregate outgoing bytes for the traffic associated with this session.

flow-reporting-fields

Specifies the flow fields and their order based on which messages should be

published.

application-id

Reports the application/category ID that is classified for this flow.

destination-ip

Reports the destination IP address of the traffic.

destination-transport-port

Reports the destination port of the traffic.

downlink-volume

Reports the total number of bytes received for this flow by the subscriber.

flow-end-milli-seconds

Reports the timestamp (milli-seconds) in UNIX time format when the flow ends.

flow-end-seconds

Reports the timestamp (seconds) in UNIX time format when the flow ends.

flow-start-milli-seconds

Reports the timestamp (milli-seconds) in UNIX time format when the flow starts.

flow-start-seconds

Reports the timestamp (seconds) in UNIX time format when the flow starts.

observation-time-seconds

Reports the timestamp (seconds) of the record.

protocol-identifier

Reports the transport layer protocol of the flow (TCP or UDP).

record-type

Reports the reporting record type of the flow: 0 - flow start, 1 - flow end, 2 - flow interim.

report-id

Reports the reporting module ID.

report-version

Reports the format version of this record.

route-domain

Reports the route domain ID of the flow.

source-ip

Reports the source IP address of the subscriber that initiates the flow.

source-transport-port

Reports the source port of the subscriber.

subscriber-id

Reports the subscriber ID that initiates this flow.

subscriber-id-type

Reports the ID type of the subscriber that initiates this flow.

timestamp-msec

Reports the timestamp (milli-seconds) of the record.

total-transactions

Reports the total number of transactions of this flow.

uplink-volume

Reports the number of bytes sent from the subscriber in this flow.

url-category-id

Reports the ID of the first URL category that is classified for the flow.

vlan-id

Reports the Vlan ID of the flow.

transaction-reporting-fields

Specifies the transaction fields and their order based on which messages should be published.

application-id

Reports the application/category ID that is classified for this transaction.

destination-ip

Reports the destination IP address of the traffic.

destination-transport-port

Reports the destination port of the traffic.

`downlink-volume`
Reports the number of HTTP response bytes for this transaction.

`http-hostname`
Reports the HTTP host name of this traffic.

`http-hostname-truncated`
Reports the truncated HTTP host name due to excessive length.

`http-response-code`
Reports the HTTP response code of the transaction.

`http-url`
Reports the HTTP URL of the transaction.

`http-url-truncated`
Reports the truncated HTTP URL of the transaction due to excessive length.

`http-user-agent`
Reports the user agent of the HTTP request in this transaction.

`http-user-agent-truncated`
Reports the truncated user agent of the HTTP request in this transaction due to excessive length.

`protocol-identifier`
Reports the transport layer protocol of the traffic (TCP or UDP).

`record-type`
Reports the reporting record type as 10-transactional.

`report-id`
Reports the reporting module ID.

`report-version`
Reports the format version of the transaction record.

`route-domain`
Reports the route domain ID of the traffic.

`skipped-transactions`
Reports the number of transactional reports skipped within the flow since the last successfully transmission in the transaction.

`source-ip`
Reports the source IP address of the subscriber.

`source-transport-port`
Reports the source port of the subscriber.

`subscriber-id`
Reports the subscriber ID that initiates this transaction.

`subscriber-id-type`
Reports the subscriber ID type of the subscriber that initiates this transaction.

`transaction-classification-result`
Reports all the classification tokens from the classification engine.

`transaction-end-milli-seconds`
Reports the transaction timestamp (milli-seconds) in UNIX time format when the corresponding HTTP response is received.

`transaction-end-seconds`
Reports the transaction timestamp (seconds) in UNIX time format when the corresponding HTTP response is received.

`transaction-number`
Reports the sequential number of transaction in this flow (starting from 1).

`transaction-start-milli-seconds`
Reports the transaction timestamp (milli-seconds) in UNIX time format when an HTTP request is received.

`transaction-start-seconds`
Reports the transaction timestamp (seconds) in UNIX time format when an HTTP request is received.

`uplink-volume`
Reports the number of HTTP request bytes for this transaction.

`url-category-id`
Reports the ID of the first URL category that is classified for the transaction.

vlan-id

Reports the Vlan ID of traffic.

radius-accounting

You can configure the following options for radius-accounting endpoint.

radius-aaa-virtual

Specifies the internal virtual server for radius-accounting endpoint.

sd You can configure the following options for sd endpoint.

application-reporting

Specifies whether the application reporting is enabled. When it is enabled, the APPLICATION_START and APPLICATION_STOP Event-Triggers will be reported when the application start/stop is detected. The default value is disabled.

monitoring-key

Specifies the monitoring-key.

granularity

Specifies the type of reporting will be generated when the policy applies. The options are flow, session and transaction. The default value is session which indicates the session report will be generated if this policy applies.

interval

Specifies the time interval in seconds the report will be generated. The default value is 0 which indicates this feature is disabled.

transaction

You can configure the following options when the transaction report granularity is selected.

http Specifies the HTTP transaction report options for the following HTTP attributes.

hostname-len

Specifies the maximum HTTP hostname string length to include in the HTTP transaction report. The range is 0 to 65535. The default value is 0.

uri-len

Specifies the maximum HTTP URI string length to include in the HTTP transaction report. The range is 0 to 65535. The default value is 256.

user-agent-max

Specifies the maximum HTTP user agent string length to include in the HTTP transaction report. The range is 0 to 65535. The default value is 0.

volume

You can configure the following options for volume threshold. The report will be generated when any of the following conditions happened. If reporting dest is set, either interval must be set to non-0 or one of volume properties must be set to non-0.

downlink

The report will be generated if the downlink traffic exceeds the threshold. The default value is 0 which indicates this feature is disabled.

total

The report will be generated if the uplink and downlink traffic exceeds the threshold. The default value is 0 which indicates this feature is disabled.

uplink

The report will be generated if the uplink traffic exceeds the threshold. The default value is 0 which indicates this feature is disabled.

quota

You can configure the following options for quota management.

rating-group

Specifies the rating-group name.

reporting-level

Specifies the quota reporting level whether per rating group or per service-id.

qos-rate-pir-downlink

Specifies the configured bandwidth control policy for Peak Information Rate (PIR) to apply to downlink traffic that matches this rule. Use none to reset bwc policy name or category name.

qos-rate-pir-uplink

Specifies the configured bandwidth control policy for Peak Information Rate (PIR) to apply to uplink traffic that matches this rule. Use none to reset bwc policy name or category name.

service-chain

Specifies where to forward the traffic affected by this rule.

sfc-action The following options can be configured for sfc-action.

path-name

Specifies the path name used by Service Function Chain (SFC) to program the path-id.

metadata-template
Specifies the SFC (Service-Function-Chain) metadata template.

tcl-filter
Specifies the tcl expression which uses iRule commands to filter the packet. It is a match if tcl-filter returns TRUE/1 or nomatch if FALSE/0. All configured filters (flow-info-filters, classification-filters, and tcl-filter) must match before rule actions are applied.

url-categorization-filters
Adds, deletes, or replaces a set of url-categorization-filters. You can configure the following options for a url-categorization-filter.

app-service
Specifies the name of the application service to which the url-categorization-filter belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the rule. Only the application service can modify or delete the url-categorization-filter.

url-category
Specifies the name of the url-category of the traffic where the rule applies. The default value is none.

operation
The options match and nomatch indicate the traffic flow must match or not match the condition specified in the classification filter. The default value is match.

status
Specifies the current status of the policy. The options are disabled and enabled. The default value is enabled.

SEE ALSO

create, delete, edit, glob, list, ltm profile qoe, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2018-11-01 pem policy(1)

pem profile diameter-endpoint

NAME
diameter-endpoint - Configures a Diameter endpoint profile.

MODULE
pem profile

SYNTAX
Configures the diameter-endpoint profile within the pem profile module using the syntax shown in the following sections.

MODIFY
create diameter-endpoint [name]
modify diameter-endpoint [name]
options:
 defaults-from [[name] | none]
 destination-host [string]
 destination-realm [string]
 fatal-grace-time {
 options:
 }
 enabled [yes | no]
 time [integer]
 }
 msg-max-retransmits [integer]
 msg-retransmit-delay [integer]
 origin-host [string]
 origin-realm [string]
 pem-protocol-profile-gx [[profile_name] | none]
 pem-protocol-profile-gy [[profile_name] | none]
 product-name [string]
 supported-apps [Gx | Gy | Sd]
 gx-session-id-prefix [string]

```
gy-session-id-prefix [string]
sd-session-id-prefix [string]
}
```

```
edit diameter-endpoint [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

```
reset-stats diameter-endpoint
```

```
reset-stats diameter-endpoint [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list diameter-endpoint
```

```
list diameter-endpoint [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config diameter-endpoint
```

```
show running-config diameter-endpoint [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
  one-line
  partition
```

```
show diameter-endpoint
```

```
show diameter-endpoint [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global
```

DESCRIPTION

You can use the diameter-endpoint component to modify or display a diameter-endpoint profile.

EXAMPLES

```
create my_diam_endpt defaults-from diameter-endpoint add { Gx }
```

Creates a custom diameter-endpoint profile name my_diam_endpt that inherits its settings from the system default diameter-endpoint profile with Gx capability.

```
modify my_diam_endpt origin-host example-host.example-realm.org origin-realm example-realm destination-host example-peer.peer-realm.org destination-realm peer-realm.org
```

Modifies the origin-host and destination-host of a diameter-endpoint profile named my_diam_endpt.

```
modify my_diam_endpt msg-max-retransmits 8 msg-retransmit-delay 10000
```

Modifies the maximum times a message will be retransmitted to 8 and the retransmission delay to 10 seconds of a diameter-endpoint profile named my_diam_endpt.

OPTIONS

defaults-from

Specifies the name of the object to inherit the settings from.

destination-host

Specifies the destination host for diameter messages. This should be a FQDN.

destination-realm

Specifies the destination realm for diameter messages. This should be a FQDN.

fatal-grace-time

You can configure following options for fatal-grace-time. It defines the period that a diameter connection can be down before all sessions associated with that diameter endpoint are terminated. If the connection is re-established before fatal-grace-time seconds then the sessions will not be terminated automatically.

enabled

Specifies whether fatal-grace-time option is enabled or no.

time

Specifies the fatal-grace-time period in seconds.

msg-max-retransmits

Specifies the number of times an outgoing request message will be retransmitted before being dropped.

msg-retransmit-delay

Specifies the delay in milliseconds after which an unanswered request will be retransmitted.

origin-host

Specifies the origin host for diameter messages. This should be a FQDN.

origin-realm

Specifies the origin realm for diameter messages. This should be a FQDN.

pem-protocol-profile-gx

Specifies PEM protocol profile to be used when subscriber discovery is enabled. This protocol profile defines mapping of Diameter AVPs to subscriber ID and other PEM subscriber session attributes for

Diameter application.

pem-protocol-profile-gy

Specifies PEM protocol profile to be used when subscriber discovery is enabled. This protocol profile defines mapping of Diameter AVPs to subscriber ID and other PEM subscriber session attributes for Diameter application.

product-name

Specifies the string used in the product-name AVP in the capabilities exchange messages.

supported-apps

Adds, deletes, or replaces a set of the supported applications. Gx and Sd are mutually exclusive and cannot be added at the same time. When Gx or Sd is added, pem-protocol-profile-gx can be specified for the protocol profile. When Gy is added, pem-protocol-profile-gy can be specified for the protocol profile.

gx-session-id-prefix

Specifies the string used as Gx session id prefix in Gx protocol messages exchanged between PEM and PCRF.

gy-session-id-prefix

Specifies the string used as Gy session id prefix in Gy protocol messages exchanged between PEM and OCS.

sd-session-id-prefix

Specifies the string used as Sd session id prefix in Sd protocol messages exchanged between PEM and PCRF.

SEE ALSO

edit, glob, list, ltm virtual, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem policy, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016, 2017. All rights reserved.

BIG-IP 2017-09-19 pem profile diameter-endpoint(1)

pem profile radius-aaa

NAME

radius-aaa - Configures a PEM radius AAA profile.

MODULE

pem profile

SYNTAX

Configures the radius-aaa profile within the pem profile module using the syntax shown in the following sections.

CREATE/MODIFY

create radius-aaa [name]

modify radius-aaa [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

retransmission-timeout [integer]

shared-secret [string]

password [string]

transaction-timeout [integer]

edit radius-aaa [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list radius-aaa

list radius-aaa [[name] | [glob] | [regex]] ...]

show running-config radius-aaa

show running-config radius-aaa [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

```
show radius-aaa
show radius-aaa [ [ [name] | [glob] | [regex] ] ... ]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
  global
```

```
DELETE
delete radius-aaa [name]
```

DESCRIPTION

You can use the radius-aaa component to create, modify, display, or delete a radius-aaa profile.

EXAMPLES

```
create radius-aaa my_radius_aaa_profile
```

Creates a custom radius-aaa profile named my_radius_aaa_profile.

```
list radius-aaa my_radius_aaa_profile
```

Displays the properties of the radius-aaa profile named my_radius_aaa_profile.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is radiusaaa.

description

User defined description.

retransmission-timeout

Specifies the retransmission timeout value of the Radius-AAA profile in seconds.

shared-secret

Specifies the shared secret of the Radius-AAA profile when connecting to the RADIUS server.

password

Specifies the password of the Radius-AAA profile for authenticating to the RADIUS server.

transaction-timeout

Specifies the transaction timeout value of the Radius-AAA profile in seconds.

partition

Specifies the administrative partition within which the profile resides.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-04-21 pem profile radius-aaa(1)

pem profile spm

NAME

spm - Configures a Subscriber Policy Manager profile.

MODULE

pem profile

SYNTAX

Configures the spm profile within the pem profile module using the syntax shown in the following sections.

CREATE/MODIFY

```
create spm [name]
```

```
modify spm [name]
```

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

global-policies-high-precedence [add | delete | replace-all-with] {
[policy_name ...]

```

}
global-policies-high-precedence [ default | none ]
global-policies-low-precedence [add | delete | replace-all-with] {
  [policy_name ...]
}
global-policies-low-precedence [ default | none ]
unknown-subscriber-policies [add | delete | replace-all-with] {
  [policy_name ...]
}
unknown-subscriber-policies [ default | none ]
fast-pem [enable | disable]
fast-vs-name [name]

```

```
edit spm [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```

all-properties
non-default-properties

```

```
reset-stats spm
```

```
reset-stats spm [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list spm
```

```
list spm [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config spm
```

```
show running-config spm [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```

all-properties
non-default-properties
one-line
partition

```

DELETE

```
delete spm [name]
```

DESCRIPTION

You can use the spm component to create, modify, display, or delete a spm profile.

EXAMPLES

```
create spm my_spm_profile
```

Creates a custom spm profile named my_spm_profile.

```
list spm my_spm_profile
```

Displays the properties of the spm profile named my_spm_profile.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile. The default value is spm.

description

User defined description.

global-policies-high-precedence

Adds, deletes, or replaces a set of the policies.

global-policies-low-precedence

Adds, deletes, or replaces a set of the policies.

unknown-subscriber-policies

Adds, deletes, or replaces a set of the policies.

fast-pem

Specifies whether fast PEM optimization is enabled or not. The default is enabled.

Specifies whether fast PEM optimization is enabled or not. PEM optimization will use the fast-vs-name virtual server for a portion of the traffic. The default is enabled.

fast-vs-name

Specifies the virtual server which will be used in fast PEM optimization when fast-pem is enabled. The virtual server should have fastL4 profile attached.

partition

Specifies the administrative partition within which the profile resides.

SEE ALSO

edit, glob, list, ltm virtual, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem policy, pem profile diameter-endpoint, pem reporting format-script, pem service-chain-endpoint, pem subscriber, regex, reset-stats, tmssh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2017-03-06 pem profile spm(1)

pem profile subscriber-mgmt

NAME

subscriber-mgmt - Configures a Subscriber Management profile.

MODULE

pem profile

SYNTAX

Configures the subscriber-mgmt profile within the pem profile module using the syntax shown in the following sections.

CREATE/MODIFY

create subscriber-mgmt [name]

modify subscriber-mgmt [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

dhcp-lease-query {

options:

enabled [false | true]

vs-name [name]

}

drop-unknown-traffic-from-server-side [enabled | disabled]

sess-creation-from-server-side [enabled | disabled]

edit subscriber-mgmt [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list subscriber-mgmt

list subscriber-mgmt [[[name] | [glob] | [regex]] ...]

show running-config subscriber-mgmt

show running-config subscriber-mgmt [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

app-service

description

dhcp-lease-query {

options:

enabled

vs-name

}

drop-unknown-traffic-from-server-side

sess-creation-from-server-side

DELETE

delete subscriber-mgmt [name]

DESCRIPTION

You can use the subscriber-mgmt component to create, modify, display, or delete a subscriber management profile.

EXAMPLES

```
create subscriber-mgmt my_subscriber_mgmt_profile
```

Creates a custom subscriber management profile named my_subscriber_mgmt_profile.

```
list subscriber-mgmt my_subscriber_mgmt_profile
```

Displays the properties of the subscriber management profile named my_subscriber_mgmt_profile.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings

and values from the parent profile. The default value is subscriber-mgmt.

description
User defined description.

dhcp-lease-query
Specifies the DHCP lease query attributes to be used for traffic triggered subscribers.

The options are:

enabled
The default is disabled.

vs-name
Specifies the name of the DHCP virtual server that is used to handle lease query and reply.

sess-creation-from-server-side
Specifies whether session creation from server side is enabled or not. The default is enabled.

drop-unknown-traffic-from-server-side
Specifies whether drop unknown traffic from server side is enabled or not. The default is disabled.

partition
Specifies the administrative partition within which the profile resides.

SEE ALSO

edit, glob, list, ltm virtual, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, pem subscriber, regex, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2018-08-22 pem profile subscriber-mgmt(1)

pem protocol diameter-avp

NAME
diameter-avp - Configures diameter AVPs in Policy Enforcement Manager (PEM).

MODULE
pem protocol

SYNTAX
Configure the diameter-avp component within the pem protocol module using the syntax shown in the following sections.

CREATE/EDIT/MODIFY
create diameter-avp [name]
modify diameter-avp [name]
options:
app-service [[string] | none]
avp-code [integer]
data-type [address | enumerated | float32 | float64 | grouped | integer32 | integer64 | octetstring | rat-type | time | unsigned32 | unsigned64]
description [string]
length [integer]
parent-avp [[diameter_avp_name | none]
vendor-id [integer]

edit diameter-avp [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY
list diameter-avp
list diameter-avp [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line
partition

DELETE
delete diameter-avp [name]

Note: You must remove all references to a diameter-avp before you can delete the diameter-avp.

DESCRIPTION

You can use the diameter-avp component to configure Diameter AVP definitions in Policy Enforcement Manager.

EXAMPLES

```
create diameter-avp user_equipment_value { data-type octetstring avp-code 460 }
```

Creates a PEM diameter-avp user_equipment_value with avp-code 460 and data-type octetstring.

```
delete diameter-avp user_equipment_value
```

Deletes the diameter-avp named user_equipment_value.

```
list diameter-avp user_equipment_value
```

Displays the properties of the diameter-avp named user_equipment_value.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

avp-code

Specifies the avp-code of the Diameter AVP.

data-type

Specifies the data type of the Diameter AVP. The default value is octetstring.

Note: The 'V' bit, known as the Vendor-Specific bit, indicates whether the optional vendor-id field is present in the AVP header. When set, the AVP code belongs to the specific vendor code address space. The 'V' bit is defined as flag-vendor-specific property in pem protocol profile gx.

The options are:

address

The address format is derived from the octetstring AVP base format. It is a discriminated union, representing, for example a 32-bit (IPv4) or 128-bit (IPv6) address, most significant octet first. The first two octets of the address AVP represents the AddressType, which contains in address family. The AddressType is used to discriminate the content and format the remaining octets.

enumerated

Enumerated is derived from the integer32 AVP Base Format. The definition contains a list of valid values and their interpretation and is described in the Diameter application introducing the AVP.

float32

This represents floating point values of single precision. The 32-bit value is transmitted in network byte order. The AVP length field MUST be set to 12 (16 if the 'V' bit is enabled).

float64

This represents floating point values of double precision. The 64-bit value is transmitted in network byte order. The AVP length field MUST be set to 16 (20 if the 'V' bit is enabled).

grouped

The data field is specified as a sequence of AVPs. Each of these AVPs follows - in the order in which they are specified - including their headers and padding. The AVP length field is set to 8 (12 if the 'V' bit is enabled) plus the total length of all included AVPs, including their headers and padding. Thus the AVP length field of an AVP of type grouped is always a multiple of 4.

integer32

32 bit signed value, in network byte order. The AVP length field MUST be set to 12 (16 if the 'V' bit is enabled).

integer64

64 bit signed value, in network byte order. The AVP length field MUST be set to 16 (20 if the 'V' bit is enabled).

octetstring

The data contains arbitrary data of variable length. Unless otherwise noted, the AVP length field MUST be set to at least 8 (12 if the 'V' bit is enabled). AVP Values of this type that are not a multiple of four-octets in length is followed by the necessary padding so that the next AVP (if any) will start on 32-bit boundary.

rat-type

specifies the value format to be encoded or decoded as the RAT-Type defined in 3GPP TS 29.212.

time The time format is derived from the octetstring AVP base format. The string MUST contain four octets, in the same format as the first four bytes are in the NTP timestamp format.

unsigned32

32 bit unsigned value, in network byte order. The AVP length field MUST be set to 12 (16 if the 'V' bit is enabled).

unsigned64

64 bit signed value, in network byte order. The AVP length field MUST be set to 16 (20 if the 'V' bit is enabled).

bit is enabled).

utf8string

The utf8string format is derived from the octetstring AVP base format. This is a human readable string represented using the ISO/IEC IS 10646-1 character set, encoded as an octetstring using the UTF-8 transformation format described in RFC 2279.

description
User defined description.

length
Specifies the data length of the Diameter AVP.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

parent-avp
Specifies the name of the parent AVP if it is in a grouped AVP.

vendor-id
Specifies the vendor-id of the Diameter VSA.

SEE ALSO

create, delete, edit, glob, list, modify, pem protocol profile gx, pem protocol profile radius, pem protocol radius-avp, pem subscriber-attribute, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-03-14 pem protocol diameter-avp(1)

pem protocol profile gx

NAME

diameter - Configures a Diameter protocol profile in Policy Enforcement Manager (PEM).

MODULE

pem protocol profile

SYNTAX

Configure the diameter component within the pem protocol profile module using the syntax shown in the following sections.

CP/CREATE/EDIT/MODIFY

```
cp diameter [source_name] [destination_name]
```

```
create diameter [name]
```

```
modify diameter [name]
```

options:

```
app-service [[string] | none]
```

```
description [string]
```

```
messages [add | delete | modify | replace-all-with] {  
  [ [message-name] ] {
```

options:

```
direction [any | in | out]
```

```
message-type [ccr-i | cca-i | ccr-u | cca-u | ccr-t | rar | raa]
```

```
avps [add | delete | modify | replace-all-with] {
```

```
  [ [avp-name] ] {
```

options:

```
default [string]
```

```
diameter-avp [ [diameter_avp_name] | none]
```

```
flag-mandatory [disabled | enabled]
```

```
flag-protected [disabled | enabled]
```

```
flag-vendor-specific [disabled | enabled]
```

```
interim-message-include [disabled | enabled]
```

```
parent-label [string]
```

```
reporting-message-include [disabled | enabled]
```

```
subscriber-attr [ [subscriber_attribute_name] | none]
```

```
}
```

```
}
```

```
}
```

```
}
```

```
subscriber-id {
```

```
  avp [ [diameter_avp_name] | none]
```

```
  type [e164 | imsi | nai | private]
```

```
  type-avp [ [diameter_avp_name] | none]
```

```
}
```

```
edit diameter [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
all-properties
```

```
non-default-properties
```

```
DISPLAY
```

```
list diameter
```

```
list diameter [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config diameter
```

```
show running-config diameter [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
all-properties
```

```
non-default-properties
```

```
one-line
```

```
partition
```

```
DELETE
```

```
delete diameter [name]
```

Note: You must remove all references to a PEM protocol profile diameter before you can delete it.

DESCRIPTION

You can use the diameter component to configure PEM protocol profile diameter definitions in Policy Enforcement Manager.

EXAMPLES

```
create cust_diam messages add {my_ccr direction out message-type ccr { avps add {avp1 { subscriber-attr 3gpp_location diameter-avp user_equipment_value flag-mandatory enabled} } } }
```

Creates a custom PEM Diameter protocol profile `cust_diam` and adds a message. The message is defined as CCR on the egress direction. PEM will insert the Diameter AVP as specified in `user_equipment_value` with the value stored subscriber attribute `3gpp_location` with mandatory flag enabled.

```
delete diameter cust_diam
```

Deletes the PEM diameter protocol profile named `cust_diam`.

```
list diameter cust_diam
```

Displays the properties of the PEM diameter protocol profile named `cust_diam`.

OPTIONS

```
app-service
```

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

```
description
```

User defined description.

```
messages
```

Adds, deletes, or replaces a set of messages which specify mapping of Diameter AVPs to subscriber session attribute for specific Diameter message. If a message by the specified name does not exist, it will be created. You can configure the following options for a message:

```
app-service
```

Specifies the name of the application service to which the message belongs. The default value is none. Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the message. Only the application service can modify or delete the message.

```
direction
```

Specifies the direction of the message.

The options are:

`any` PEM will process the message in both ingress and egress directions.

`in` PEM will process the message in ingress direction.

`out` PEM will process the message in egress direction.

```
message-type
```

Specifies the type of the message.

The options are:

```
ccr-i
```

The message is Diameter Credit-Control-Request (CCR) Initial.

```
cca-i
```

The message is Diameter Credit-Control-Answer (CCA) Initial.

```
ccr-u
```

The message is Diameter Credit-Control-Request (CCR) Update.

cca-u

The message is Diameter Credit-Control-Answer (CCA) Update.

ccr-t

The message is Diameter Credit-Control-Request (CCR) Terminate.

rar The message is Diameter Re-Authorization-Request (RAR).

raa The message is Diameter Re-Authorization-Answer (RAA).

avps Adds, deletes, or replaces a set of mapping between Diameter AVPs and PEM subscriber attributes. You can configure the following options.

app-service

Specifies the name of the application service to which the AVP belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the avp. Only the application service can modify or delete the avp.

default

Specifies the Diameter AVP default value. When inserting the AVP, the default value is used if the corresponding subscriber attribute is not defined or is not present.

diameter-avp

Specifies the name of the Diameter AVP. The default value is none.

flag-mandatory

Specifies the value of the mandatory flag in the Diameter AVP when inserting into the message. This flag only applies to Diameter AVP in outgoing message.

flag-protected

Specifies the value of the protected flag in the Diameter AVP when inserting into the message. This flag only applies to Diameter AVP in outgoing message.

flag-vendor-specific

Specifies the value of the vendor-specific flag in the Diameter AVP when inserting into the message. This flag only applies to Diameter AVP in outgoing message.

interim-message-include

Specifies whether this AVP needs to be included in the interim-message (ccr-u only) updates which are generated if there is any change related to session parameters. This flag only applies to Diameter AVP in outgoing message.

parent-label

Specifies how grouped AVPs can be combined. The AVPs with the same parent-label will be combined in the same grouped AVP.

reporting-message-include

Specifies whether this AVP needs to be included in the reporting-message (ccr-u only) updates which are generated for reporting usage information. This flag only applies to Diameter AVP in outgoing message.

subscriber-attr

Specifies the name of the subscriber session attribute to be mapped to Diameter AVP. The default value is none.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

subscriber-id

Specifies how subscriber ID specified by PCRF should be read by PEM.

avp Subscription ID comes with type and data in diameter. This field specifies the avp in the message that should be matched to get subscriber ID data in raw format.

type Specifies the subscriber ID type (imsi, e164, private) that PEM will use for the session for the ID read by avp.

The options are:

e164 A numbering plan that defines the format of an MSISDN international phone number (up to 15 digits). The number typically consists of three fields: country code, national destination code, and subscriber number.

imsi International Mobile Subscriber Identity. A globally unique code number that identifies a GSM, UMTS, or LTE mobile phone user.

nai Network Access Identifier. A fully qualified network name in the form @; identifies a subscriber and the home network to which the subscriber belongs.

private

The subscriber id type is private for the given deployment.

type-avp

Subscription ID comes with type and data in diameter. type-avp specifies avp in message that should be matched.

SEE ALSO

create, delete, edit, glob, list, modify, pem protocol profile radius, pem protocol diameter-avp, pem protocol radius-avp, pem subscriber-attribute, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016-2017. All rights reserved.

BIG-IP 2017-07-19 pem protocol profile gx(1)

pem protocol profile radius

NAME

radius - Configures a RADIUS protocol profile in Policy Enforcement Manager (PEM).

MODULE

pem protocol profile

SYNTAX

Configure the radius component within the pem protocol profile module using the syntax shown in the following sections.

```
CP/CREATE/EDIT/MODIFY
cp radius [source_name] [destination_name]
create radius [name]
modify radius [name]
options:
  app-service [[string] | none]
  description [string]
  messages [add | delete | modify | replace-all-with] {
    [ [message-name] ] {
options:
  direction [any | in | out]
  message-type [acct-req-start | acct-req-stop | acct-req-interim-update]
  avps [add | delete | modify | replace-all-with] {
    [ [avp-name] ] {
      options:
        default [string]
        ingress-op [ import | none]
        radius-avp [ [radius_avp_name] | none]
        subscriber-attr [ [subscriber_attribute_name] | none]
      }
    }
  }
  subscriber-id [add | delete | modify | replace-all-with] {
    [ [id-name] ] {
options:
  order [integer]
  prefix [[string] | none]
  radius-avp [[radius_avp_name] | none]
  suffix [[string] | none]
  }
  }
  subscriber-id-type [e164 | imsi | nai | private]
```

```
edit radius [ [ [name] | [glob] | [regex] ] ... ]
```

```
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list radius
list radius [ [ [name] | [glob] | [regex] ] ... ]
show running-config radius
show running-config radius [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete radius [name]
```

Note: You must remove all references to a pem protocol profile radius before you can delete it.

DESCRIPTION

You can use the radius component to configure pem protocol profile radius definitions in Policy Enforcement Manager.

EXAMPLES

```
create cust_acct_start messages add { my_acct_start { direction in message-type acct-req-start avps add { avp1 { subscriber-attr _sys_attr_3gpp_imeisv radius-avp _sys_radius_3gpp_imeisv ingress-op import } } } }
```

Creates a custom PEM RADIUS protocol profile `cust_acct_start` and add a message to define how the RADIUS message can be processed. The message is defined as RADIUS accounting on the ingress direction. The mapping action `ingress-op` is to extract RADIUS AVP defined in `_sys_radius_3gpp_imeisv` and store the value into subscriber attribute `_sys_attr_3gpp_imeisv`.

```
delete radius cust_acct_start
```

Deletes the PEM RADIUS protocol profile named `cust_acct_start`.

```
list radius cust_acct_start
```

Displays the properties of the PEM RADIUS protocol profile named `cust_acct_start`.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

messages

Adds, deletes, or replaces a set of messages which specify mapping of RADIUS AVPs to subscriber session attributes for specific Gx message. If a message by the specified name does not exist, it will be created. You can configure the following options for a message:

app-service

Specifies the name of the application service to which the message belongs. The default value is none. Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the message. Only the application service can modify or delete the message.

direction

Specifies the direction of the message.

The options are:

`any` PEM will process the message in both ingress and egress directions.

`in` PEM will process the message in ingress direction.

`out` PEM will process the message in egress direction.

message-type

Specifies the type of the message.

The options are:

acct-req-start

The message is RADIUS accounting with the value of `Acct-Status-Type` AVP set to 1 (Start).

acct-req-stop

The message is RADIUS accounting with the value of `Acct-Status-Type` AVP set to 2 (Stop).

acct-req-interim-update

The message is RADIUS accounting with the value of `Acct-Status-Type` AVP set to 3 (Interim-Update).

`avps` Adds, deletes, or replaces a set of mapping between RADIUS AVPs and PEM subscriber attributes. You can configure the following options.

app-service

Specifies the name of the application service to which the avp belongs. The default value is none. Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the message. Only the application service can modify or delete the avp.

default

Specifies the RADIUS AVP default value. When inserting the AVP, the default value is used if the corresponding subscriber session attribute is not defined or is not present.

ingress-op

Specifies the ingress operation applied when processing the RADIUS AVP. The default value is none.

The options are:

import

Specifies that the RADIUS AVP will be parsed and the value will be stored in the subscriber attribute.

none Specifies that there is no ingress operation applied to the RADIUS AVP.

radius-avp

Specifies the name of the RADIUS AVP. The default value is none.

subscriber-attr

Specifies the name of the subscriber session attribute to be mapped to RADIUS AVP. The default value is none.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

subscriber-id

Adds, deletes, or replaces a set of RADIUS AVPs to form PEM subscriber ID. You can configure the following options:

order

Specifies the order of RADIUS AVPs when constructing the subscriber ID.

prefix

Specifies the prefix string when constructing subscriber ID with the value of the RADIUS AVP.

radius-avp

Specifies the value of RADIUS AVP which will be used to construct the subscriber ID.

suffix

Specifies the suffix string when constructing subscriber ID with the value of the RADIUS AVP.

subscriber-id-type

Specifies the subscriber ID type session attribute value for the session created.

The options are:

e164 A numbering plan that defines the format of an MSISDN international phone number (up to 15 digits). The number typically consists of three fields: country code, national destination code, and subscriber number.

imsi International Mobile Subscriber Identity. A globally unique code number that identifies a GSM, UMTS, or LTE mobile phone user.

nai Network Access Identifier. A fully qualified network name in the form @; identifies a subscriber and the home network to which the subscriber belongs.

private

The subscriber id type is private for the given deployment.

SEE ALSO

create, delete, edit, glob, list, modify, pem protocol profile gx, pem protocol diameter-avp, pem protocol gx-avp, pem subscriber-attribute, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-03-14 pem protocol profile radius(1)

pem protocol radius-avp

NAME

radius-avp - Configures RADIUS AVPs in Policy Enforcement Manager (PEM).

MODULE

pem protocol

SYNTAX

Configure the radius-avp component within the pem protocol module using the syntax shown in the following sections.

CREATE/EDIT/MODIFY

create radius-avp [name]
modify radius-avp [name]

options:

app-service [[string] | none]
data-type [3gpp-rat-type | 3gpp-user-location-info | integer | ipaddr | ipv6addr | ipv6prefix | octet | string | time]
description [string]
max-length [integer]
min-length [integer]
type [integer]
vendor-id [integer]
vendor-type [integer]

edit radius-avp [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list radius-avp

list radius-avp [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

DELETE

delete radius-avp [name]

Note: You must remove all references to a radius-avp before you can delete the radius-avp.

DESCRIPTION

You can use the radius-avp component to configure RADIUS AVP definitions in Policy Enforcement Manager.

EXAMPLES

```
create radius-avp imeisv { data-type string type 26 vendor-id 10415 vendor-type 20 }
```

Creates a PEM radius-avp imeisv which is an Vendor Specific Attribute of 3GPP with type value 26.

```
delete radius-avp imeisv
```

Deletes the radius-avp named imeisv.

```
list radius-avp imeisv
```

Displays the properties of the radius-avp named imeisv.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

data-type

Specifies the data type of the RADIUS AVP. The default value is string.

The options are:

3gpp-rat-type

specifies the value format to be encoded or decoded as the 3GPP-RAT-Type defined in 3GPP TS 29.061.

3gpp-user-location-info

specifies the value format to be encoded or decoded as the 3GPP-User-Location-Info defined in 3GPP TS 29.061.

integer

32-bit unsigned integer in network byte order.

ipaddr

IPv4 address in network byte order.

ipv6addr

IPv6 address in network byte order.

ipv6prefix

IPv6 prefix data format is defined in RFC 3162.

octet

UTF-8 text [RFC3629], totaling 253 octets or less in length.

string

string (i.e., binary data), totaling 253 octets or less in length. This includes the opaque encapsulation of data structures defined outside of RADIUS.

time time as a 32-bit unsigned value in network byte order and in seconds since 00:00:00 UTC, January 1, 1970.

description

User defined description.

max-length

Specifies the maximum data length of the RADIUS AVP/VSA. It doesn't include the AVP/VSP header. The default value is 253.

min-length

Specifies the minimum data length of the RADIUS AVP/VSA. It doesn't include the AVP/VSP header. The default value is 1.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

type Specifies the type of the RADIUS AVP. 26 is for vendor specific attribute (VSA).

vendor-id

Specifies the vendor-id of the RADIUS VSA. This property is mandatory if type is 26.

vendor-type

Specifies the vendor-type of the RADIUS VSA. This property is mandatory if type is 26.

SEE ALSO

create, delete, edit, glob, list, modify, pem protocol profile gx, pem protocol profile radius, pem protocol diameter-avp, pem subscriber-attribute, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-03-14 pem protocol radius-avp(1)

pem quota-mgmt rating-group

NAME

rating-group - Configures a rating-group for quota management in Policy Enforcement Manager (PEM).

MODULE

pem quota management

SYNTAX

Modify the rating-group component within the pem quota-mgmt module using the syntax shown in the following sections.

CREATE/MODIFY

create rating-group [name]

modify rating-group [name]

options:

app-service [[string] | none]

rating-group-id [integer]

description [string]

request-on-install [yes | no]

default-threshold [integer]

default-validity-time [integer]

default-quota-holding-time [integer]

initial-quota-request {

interval [integer]

volume {

input-octets

output-octets

total-octets

}

default-quota {

interval [integer]

volume {

input-octets

output-octets

total-octets

}

time {

usage-time

consumption-time

}

default-breach-action [terminate | allow | redirect]

default-forwarding-endpoint [name]

edit rating-group [[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

```
list rating-group
list rating-group [ [ [name] | [glob] | [regex] ] ... ]
show running-config rating-group
show running-config rating-group [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
```

DELETE

```
delete rating-group [name]
```

Note: You must remove all references to a rating-group object before you can delete it.

EXAMPLES

```
create rating-group rg1 {
  rating-group-id 1
  initial-quota-request {
    volume {
      input-octets 1000  output-octets 1000  total-octets 2000
    }
  }
  default-quota {
    volume {
      input-octets 1000  output-octets 1000  total-octets 2000
    }
  }
  request-on-install yes }
```

Creates a PEM rating-group named rg1.

```
delete rating-group rg1
```

Deletes the rating-group named rg1.

```
list rating-group rg1
```

Displays the properties of the rating-group named rg1.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

rating-group-id

Specifies the rating-group-id that will be used by quota managing endpoint. For instance, this could be the rating group in case of Gy endpoint.

request-on-install

Specifies whether quota has to be requested from the quota managing endpoint (Eg : Gy) when policy referring to this rating-group is installed for a subscriber or later when flow is initiated.

default-threshold

Specifies the default threshold if the quota managing endpoint does not specify threshold.

default-validity-time

Specifies the default validity time for the quota in seconds if OCS did not specify it.

default-quota-holding-time

Specifies the default quota holding time in seconds for which quota is valid without any usage if not specified by OCS.

initial-quota-request

Specifies the initial quota, that will be requested from the quota managing endpoint. Could be either time or volume.

time Specifies the time in seconds.

volume

You can configure the following options for volume initial quota.

output-octets

Specifies the initial quota for downlink traffic.

total-octets

Specifies the initial quota for total uplink and downlink traffic.

input-octets

Specifies the initial quota for uplink traffic.

default-quota

Specifies the default quota, that will be used if quota managing endpoint does not respond. Could be either time or volume.

time Specifies the quota in time.

usage-time

Specifies the usage time in seconds.

consumption-time

Specifies the quota consumption time in seconds.

volume

You can configure the following options for volume default quota.

output-octets

Specifies the default quota for downlink traffic.

total-octets

Specifies the default quota for total uplink and downlink traffic.

input-octets

Specifies the default quota for uplink traffic.

SEE ALSO

create, delete, edit, glob, list, modify, pem policy, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2016-01-07 pem quota-mgmt rating-group(1)

pem reporting format-script

NAME

format-script - Configures format scripts for the Policy Enforcement Manager (PEM).

MODULE

pem reporting

SYNTAX

Modify the format-script component within the pem reporting module using the syntax shown in the following sections.

CREATE/MODIFY

create format-script [name]

modify format-script [name]

options:

app-service [[string] | none]

definition [string]

description [string]

edit format-script [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list format-script

list format-script [[[name] | [glob] | [regex]] ...]

show running-config format-script

show running-config format-script [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete format-script [name]

Note: You must remove all references to a format script object before you can delete it.

DESCRIPTION

You can use the format-script component to create scripts for HSL reporting. The scripts use TCL syntax and define a custom format that is applied in an enforcement policy rule. The format and fields available differ

depending on whether the rule specifies session-based or flow-based reporting.

EXAMPLES

```
create format-script fm1 { definition { return "(flow app_id[PEM::flow stats reported app-id], bytes-in:[PEM::flow stats reported bytes-in])" } }
```

Creates a PEM reporting format script named fm1.

```
delete format-script fm1
```

Deletes the format script named fm1.

```
list format-script fm1
```

Displays the properties of the format script named fm1.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

definition

Specifies a script using TCL syntax that defines a custom format for HSL reporting applied in an enforcement policy rule. The format and fields available differ depending on whether you are using session-based or flow-based reporting in the rule.

Session-based formats:

The options are:

app-id

Specifies the application id of the session record.

bytes-in

Specifies the aggregate incoming bytes of the session.

bytes-out

Specifies the aggregate outgoing bytes of the session.

calling-station-id

Specifies the calling station Id.

called-station-id

Specifies the called station Id.

flows-concurrent

Specifies the number of concurrent flows in the session.

flows-duration

Specifies the duration of the flow in the session.

flows-ended

Specifies the number of flows ended in the session.

flows-new

Specifies the number of new flows in the session.

last-sent-sec

Specifies the value of seconds of the timestamp since the previous record was sent.

param-3gpp

Specifies the comma-separated string of the value of imsi, imeisv, tower-id, and user-name.

record-reason

Specifies the reason for sending report. The values are 1: period time, 2: volume threshold, 3: subscriber logout, 4: inactivity.

record-type

Specifies the type of the session-based record (always 3).

report-id

Specifies the report Id.

report-version

Specifies the report version.

subscriber-id

Specifies the subscriber id.

subscriber-id-type

Specifies the subscriber id type (e164, imsi, nai, or private).

timestamp

Specifies the seconds value of the timestamp when the record was generated. The Unix epoch is 1970-01-01T00:00:00Z.

timestamp-milliseconds

Specifies the milliseconds value of the timestamp when the record was generated.

Note: Recommend not to use the following options. They are available for backward compatibility.

last-send-usec

Specifies the value of microseconds of the timestamp since the previous record was sent.

rec-reason

Specifies the reason for sending report. The values are 1: period time, 2: volume threshold, 3: subscriber logout, 4: inactivity.

rec-type

Specifies the type of the session-based record (always 3).

subs-id

Specifies the subscriber id.

subs-id-type

Specifies the subscriber id type (e164, imsi, nai, or private).

timestamp-sec

Specifies the seconds value of the timestamp when the record was generated. The Unix epoch is 1970-01-01T00:00:00Z.

timestamp-usec

Specifies the microseconds value of the timestamp when the record was generated.

Flow-based formats:

The options are:

app-id

Specifies the application id of the flow record.

bytes-in

Specifies the aggregate incoming bytes of the flow.

bytes-out

Specifies the aggregate outgoing bytes of the flow.

dst-ip

Specifies the destination ip address of the flow.

dst-port

Specifies the destination port of the flow.

flow-start-milli-seconds

Specifies the milliseconds value of the timestamp when the flow starts.

flow-start-time-sec

Specifies the seconds value of the timestamp when the flow starts. The Unix epoch is 1970-01-01T00:00:00Z.

flow-end-milli-seconds

Specifies milliseconds value of the timestamp when the flow ends.

flow-end-time-sec

Specifies the seconds value of the timestamp when the flow ends. The Unix epoch is 1970-01-01T00:00:00Z.

proto

Specifies the protocol of the flow.

record-type

Specifies the type of the flow-based record. The value is 0: flow init, 1: flow interim, and 2: flow end.

report-id

Specifies the report Id.

route-domain

Specifies the route domain Id of the traffic.

report-version

Specifies the report version.

src-ip

Specifies the source ip address of the flow.

src-port

Specifies the destination port of the flow.

subscriber-id

Specifies the subscriber id.

subscriber-id-type

Specifies the subscriber id type (e164, imsi, nai, or private).

timestamp

Specifies the of seconds value of the timestamp when the record was generated. The Unix epoch is 1970-01-01T00:00:00Z.

timestamp-milliseconds

Specifies the milliseconds value of the timestamp when the record was generated.

transactions

Specifies the number of transactions in the flow.

urlcat-id

Specifies the unique number that represents the first (most relevant) URL category that is classified for the flow.

vlan-id

Specifies the vlan Id of the traffic.

Note:Recommend not to use the following options. They are available for backward compatibility.

flow-start-time-usec

Specifies microseconds value of the timestamp when the flow starts.

flow-end-time-usec

Specifies microseconds value of the timestamp when the flow ends.

rec-type

Specifies the type of the flow-based record. The value is 0: flow init, 1: flow interim, and 2: flow end.

subs-id

Specifies the subscriber id.

subs-id-type

Specifies the subscriber id type (e164, imsi, nai, or private).

timestamp-sec

Specifies the of seconds value of the timestamp when the record was generated. The Unix epoch is 1970-01-01T00:00:00Z.

timestamp-usec

Specifies the microseconds value of the timestamp when the record was generated.

description

Specifies a user-defined description.

SEE ALSO

create, delete, edit, glob, list, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem service-chain-endpoint, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2017-03-02 pem reporting format-script(1)

pem service-chain-endpoint

NAME

service-chain-endpoint - Configures service chain endpoints for the Policy Enforcement Manager (PEM).

MODULE

pem

SYNTAX

Modify the service-chain-endpoint component within the pem module using the syntax shown in the following sections.

CREATE/MODIFY

```
create service-chain-endpoint [name]
modify service-chain-endpoint [name]
options:
  app-service [[string] | none]
  service-endpoints [add | delete | modify | replace-all-with] {
[service endpoint name ... ] {
  options:
```

```

app-service [[string] | none]
forwarding-endpoint
  to-endpoint [forwarding endpoint name]
from-vlan [vlan name]
http-adapt-service
  internal-virtual [internal virtual server | none]
  icap-type [request | response | both | none]
order [integer]
service-option [optional | mandatory]
steering-policy [policy name | none]
}
}

```

```
edit service-chain-endpoint [ [name] | [glob] | [regex] ] ... ]
```

```

options:
  all-properties
  non-default-properties

```

DISPLAY

```

list service-chain-endpoint
list service-chain-endpoint [ [name] | [glob] | [regex] ] ... ]
show running-config service-chain-endpoint
show running-config service-chain-endpoint [ [name] | [glob] | [regex] ] ... ]

```

```

options:
  all-properties
  non-default-properties
  one-line
  partition

```

DELETE

```
delete service-chain-endpoint [name]
```

Note: You must remove all references to a service-chain-endpoint before you can delete the service-chain-endpoint.

DESCRIPTION

You can use the service-chain-endpoint component to configure service-chain-endpoint definitions for the Policy Enforcement Manager (PEM). Each service-chain-endpoint consists of one or more service-endpoints, where a service-endpoint consists of a non-zero integer order, existing from-vlan a valid fwd-endpoint or a http-adaptation-service endpoint. When you configure a BIG-IP that has a service-chain-endpoint with multiple service-endpoints, traffic will pass through different endpoints chosen dynamically.

Note: You must create a valid forwarding-endpoint and a valid vlan before you can create a service-endpoint. If you are enabling http-adapt-service, you must create Request Adapt and Response Adapt profiles and attach to the traffic virtual. Also create an internal-virtual and enable icap profile. You must also give each service-endpoint an order from 1 up to $2^{32}-1$. The lower the service-endpoint order is, the higher its precedence is (i.e., traffic will pass through it before other higher order service-endpoints). Each service-endpoint has a boolean (true/false) service-option that defines what would happen if the service-endpoint is down. If service-option is mandatory, the traffic flow is dropped if the service-endpoint is down. If service-option is optional, the traffic flow will be bypassed to the next available service-endpoint.

For more information about how to create a vlan, please refer to net vlan. Also please refer to pem forwarding-endpoint for more information about how to create a pem forwarding-endpoint.

EXAMPLES

```
create service-chain-endpoint chain1 service-endpoints add { ser_ep1 { order 10 from-vlan vlan1 forwarding-endpoint { to-endpoint fw_ep1 } service-option optional } ser_ep2 { order 5 from-vlan vlan2 http-adapt-service { internal-virtual iv1 } service-option mandatory } }
```

Creates a PEM service-chain-endpoint named chain1 that has two service-endpoints: ser_ep1 and ser_ep2. The first ser_ep1 has an order of 10 and is optional and has forwarding-endpoint with to-endpoint fw_ep1, type transparent and vlan1 as a from-vlan. The second ser_ep2 has an order of 5 is mandatory and has http-adapt-service enabled with ivs1 as internal-server and vlan2 as a from-vlan. Note that ser_ep2 will precede ser_ep1 because the lower the service-endpoint order is, the higher its precedence is.

```
delete service-chain-endpoint chain1
```

Deletes the service-chain-endpoint named chain1.

```
list service-chain-endpoint chain1
```

Displays the properties of the service-chain-endpoint named chain1.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

service-endpoints

Adds, deletes, or replaces a set of the service endpoints by specifying a series of service-endpoint names. If any of these names did not exist before, then new names will be created. Each service-endpoint is identified by a vlan and a forwarding-endpoint.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object,

you cannot modify or delete the object. Only the application service can modify or delete the object.

forwarding-endpoint
Specifies the forwarding endpoint attributes to be set. The below attributes can be set:

to-endpoint
This is a default endpoint that will be chosen if steering policy is not configured. You have to create a valid PEM forwarding-endpoint before you can add to-endpoint to a service-endpoint.

from-vlan
Specifies the vlan that the traffic will come from toward the service-endpoint. Note: The vlan has to exist before you can create a from-vlan field.

http-adapt-service
Specifies the option to set attributes for http adapt services. Below are the attributes that can be set.

internal-virtual
This is the internal virtual on which icap is enabled. You have to create the internal-virtual and assign icap profile before adding here.

icap-type
Defines the ICAP adaptation type: request only adaptation, request and response adaptation or both types of adaptations combined.

order
Specifies the order of the service-endpoint among other service-endpoints. The lower the service-endpoint's order is, the more precedence it has (i.e., the traffic will go through the lowest-ordered service-endpoint first, then through higher order service-endpoint, ... etc.).

service-option
Specifies the behavior when a service-endpoint is not available (i.e., is down). This option is limited when ICAP is defined as the service-endpoint and will not apply if the ICAP service is unavailable. You can configure the following options:

mandatory
If the service-endpoint is down, the traffic flow is dropped.

optional
If the service-endpoint is down, the traffic flow will be bypassed to the next available service-endpoint.

steering-policy
If the steering policy is configured, the policy is evaluated and if steering is enabled the flow will be steered to the corresponding endpoint.

SEE ALSO

create, delete, edit, glob, list, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem subscriber, pem subscribers, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012. All rights reserved.

BIG-IP 2016-01-07 pem service-chain-endpoint(1)

pem sessiondb

NAME
sessiondb - Displays, deletes, modifies, and reset-stats a PEM subscriber session record on the BIG-IP(r) system.

MODULE
pem

DESCRIPTION
You can use the sessiondb component to display session record on the BIG-IP system. Additionally, you can delete, reset-stats and modify a specified session record. The subscriber-id or session-ip or all and/or the filtering arguments can be specified as the query key. session-state must be specified in modify command.

When specifying an IP address in show/delete/modify/reset-stats sessiondb session-ip commands, the route domain ID may be optionally included. The route domain ID follows the IP address and is separated by '%'. For example, 10.10.10.100%5 is an IP address in route domain 5.

Note: show and reset-stats commands apply to both static and dynamic subscribers. Delete and modify commands only apply to dynamic subscribers. The session of static subscribers cannot be deleted. The session-state of static subscribers cannot be changed. To delete a static subscriber session you have to delete the static subscriber configuration in pem subscriber.

SHOW command SYNTAX

```
show sessiondb subscriber-id []
```

```
show sessiondb session-ip /[ ] []
```

```
show all
```

options:
all-properties
field-fmt

DESCRIPTION

Used to display the session information.

subscriber-id can either be the complete subscriber-id or a wildcard.

session-ip be either the complete IP address, or a subnet, by specifying the prefix.

all is used to display all the sessions. Due to existing limitations, some of the arguments cannot be used with the 'all' option. In this case, 'subscriber-id *' can be used.

all-properties options enables the per-IP traffic statistics in case of sessions with multiple IP addresses.

field-fmt option provides a structured output. This option is used for REST API compatibility.

args is an optional list of arguments, which can contain the following-

- filtering arguments: These specify the criteria for selecting sessions to display. The filtering arguments can be one or more of the following

- . imsi
- . imeisv
- . calling-station-id
- . called-station-id
- . user-name
- . session-origin
- . session-state
- . device-name
- . device-os
- . user-location-info

This can either be the user location as a hex string, or the location info broken down into individual fields as name value pairs. The values can be combined using '&' or '|' and can have only one of the operators, not both. The key value pairs allowed are as follows ('D' stands for a decimal value and 'H' for hex digit) -

- type=[CGI | SAI | RAI | TAI | TAI+ECGI] - mcc=DD[D] - mnc=DD[D] - lac=[0x]HHHH - ci=[0x]HHHH
- sac=[0x]HHHH - rac=[0x]HHHH - tac=[0x]HHHH - eci=[0x]HHHHHHH

For e.g.

```
user-location-info 0x1234567 user-location-info type=CGI&mcc=123 user-location-info  
type=TAI+ECGI|mcc=123|lac=0x3f5d
```

This is not allowed -

```
user-location-info type=CGI&mcc=123|mcc=456
```

. attr

This is the list of custom attribute key-value pairs. Multiple attributes can be combined with either '&' or '|'. (cannot have a mix of '&' and '|' operators). There can be a maximum of 30 custom attributes.

. policy

List of policy names. This can again be combined using the '&' or '|' operators. There can be maximum of 30 policies.

. blade

. tmm

- include-deleted: This specifies whether to include sessions marked for deletion in the output.
- view-mode: This specifies the format of the output and can be one of extended, table or count. The extended view mode is the default and gives a formatted output of all the session attributes and flow statistics. In the table mode, the user can select the information to be displayed by selecting the list of columns. In this mode, the 'columns' argument HAS to be specified containing the list of columns to display. Column names should be separated by ':' and should be from the following list -

. subscriber_id

. tmm_number

. blade_number

- . session_state_str
- . session_origin_str
- . ppe_session_id
- . qpe_session_id
- . imsi
- . imeisv
- . called_station_id
- . calling_station_id
- . user_name
- . device_name
- . device_os
- . tower_id
- . ip_info.ip_address
- . policy_info.policy_name
- . any custom attribute e.g. _sys_attr_3gpp_rat_type.

Any unknown column name is considered to be a custom attribute name.

If the view-mode is set to count, then only a count of the number of sessions matching the specified criteria is displayed.

- max-count: This is used to specify an upper limit on the number of matching sessions to be considered for the command.
- resume-from: This argument is used to get a page-by-page listing of the sessions, and can be used only with the field-fmt option.

DELETE command

SYNTAX

```
delete sessiondb subscriber-id []
```

```
delete session-ip [ip address>/[] []]
```

```
delete all
```

DESCRIPTION

Used to delete sessions.

The optional arguments here can specify the filtering arguments, as explained in the show command above.

All the sessions matching the criteria are deleted.

MODIFY command

SYNTAX

```
modify sessiondb subscriber-id [] session-state [marked-for-deletion | not-provisioned | provisioned | provisioning-pending ]
```

```
modify session-ip [ip address>/[] []] session-state [marked-for-deletion | not-provisioned | provisioned | provisioning-pending ]
```

DESCRIPTION

Used to modify the session state.

The optional arguments here can specify the filtering arguments, as explained in the show command above.

The state of all the matching sessions are modified to the value specified.

session-state

Specifies the new subscriber session state.

The options are:

marked-for-deletion

Specifies that the subscriber session should be scheduled for deletion.

provisioned

Specifies that the subscriber session state should be marked as provisioned, regardless of whether the policies have been assigned or not. The unknown subscriber policies are not applied to the subscriber flows, even if no subscriber policies are provisioned.

not-provisioned

Specifies that the subscriber session state should be marked as not-provisioned. No further attempts to provision the session are made. The unknown subscriber policies are applied to the subscriber flows.

provisioning-pending

Specifies that the subscriber session state should be marked as being in the process of provisioning. This will trigger a session provisioning request (e.g. Gy CCR request) immediately. If no response is received, or the provisioning process fails for any reason, another request will be sent after the retry timeout, until the session is provisioned successfully, or the number of retries is reached.

RESET-STATS command

SYNTAX

```
reset-stats sessiondb subscriber-id []
```

```
reset-stats session-ip [ip address>/[] []
```

```
reset-stats all
```

DESCRIPTION

Used to reset the session statistics.

The optional arguments here can specify the filtering arguments, as explained in the show command above.

The statistics for all the sessions matching the criteria are reset.

EXAMPLES

```
show sessiondb subscriber-id 4085551212
```

Displays the session record of subscriber id 4085551212.

```
show sessiondb subscriber-id 408*
```

Displays the session record of all subscribers with id beginning with '408'.

```
show sessiondb subscriber-id *
```

Displays the session records of all subscribers.

```
show sessiondb session-ip 10.10.10.100
```

Displays the session record of session ip address 10.10.10.100.

```
show sessiondb session-ip 10.10.10.100%5
```

Displays the session record of session ip address 10.10.10.100 in route domain 5.

```
show sessiondb session-ip 10.10.10.10/24
```

Displays the session records of session ip addresses in the 10.10.10.0 subnet.

```
show sessiondb subscriber-id 4085551212 all-properties
```

Displays the session record of subscriber id 4085551212, and statistics for each IP.

```
show sessiondb session-ip 10.10.10.100 all-properties
```

Displays the session record of session ip address 10.10.10.100, and statistics for each IP.

```
show sessiondb session-ip 10.10.10.10/24 view-mode count
```

Displays the count of the number of session records with session ip address in the 10.10.10.0 subnet.

```
show sessiondb subscriber-id 408* session-state provisioned
```

Displays the count of the number of sessions whose id begins with '408' and are in the provisioned state.

```
delete sessiondb subscriber-id 4085551212
```

Deletes the session record of subscriber id 4085551212.

```
delete sessiondb session-ip 10.10.10.100
```

Deletes the session record of IP address 10.10.10.100.

```
delete sessiondb session-ip 10.10.10.100%5
```

Deletes the session record of IP address 10.10.10.100 in route domain 5.

```
delete sessiondb session-ip session-state provisioning-pending
```

Deletes all sessions that are in the provision pending state.

```
reset-stats sessiondb subscriber-id 4085551212
```

Reset the session statistics of subscriber id 4085551212. Flows Current specifies the active flows and it cannot be reset.

modify sessiondb subscriber-id 4085551212 session-state provisioned

Modifies the session state of subscriber id 4085551212 to provisioned.

modify sessiondb subscriber-id 408* session-state provisioned

Modifies the session state of sessions with subscriber id starting with 408 to state provisioned.

SEE ALSO

delete, modify, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 pem sessiondb(1)

pem stats action

NAME

action - Displays and resets PEM policy action statistics.

MODULE

pem stats

SYNTAX

Display statistics for the action component within the pem stats module using the syntax in the following section.

DISPLAY

show action

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the action component to display PEM policy action statistics. The statistics details are described below:

Pass Specifies the number of flows that are passed (gate enabled).

Drop Specifies the number of flows that are dropped (gate disabled).

Clone

Specifies the number of flows to which clone actions apply.

HTTP Redirect

Specifies the number of flows to which redirection actions apply.

ICAP Request

Specifies the number of flows to which ICAP actions apply on the request direction.

ICAP Response

Specifies the number of flows to which ICAP actions apply on the response direction.

Steering

Specifies the number of flows to which steering actions apply.

Service Chain

Specifies the number of flows to which steering endpoint actions apply.

Steering on Response

Specifies the number of flows to which steering actions apply on the response direction.

QoS Uplink

Specifies the number of uplink flows to which QoS actions apply. Uplink means to network.

QoS Downlink

Specifies the number of downlink flows to which QoS actions apply. Downlink means to subscriber.

DSCP Marking Uplink

Specifies the number of uplink flows with DSCP action applies.

DSCP Marking Downlink

Specifies the number of downlink flows with DSCP action applies.

HTTP Headers Modify

Specifies the number of HTTP Headers Modify actions.

Insert Content

Specifies the number of Insert Content actions.

iRule

Specifies the number of iRule actions.

L2 Marking Uplink

Specifies the number of uplink flows to which L2 Marking actions apply.

L2 Marking Downlink

Specifies the number of downlink flows to which L2 Marking actions apply.

Flow Reporting

Specifies the number of actions of flow reporting record generation applied.

Session Reporting

Specifies the number of actions of session record generation applied.

Transaction Reporting

Specifies the number of actions of transaction record generation applied.

Policy Re-evaluation Rate (count/min)

Specifies the number of successful policy reevaluations per minute.

Policy Re-evaluation Rate Maximum

Specifies the maximum number of policy reevaluations overall for all subscribers and flows.

You can reset the PEM policy action statistics using reset-stats command.

EXAMPLES

show action

Displays the PEM policy action statistics.

reset-stats action

Resets the PEM policy action statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats gx, pem stats gy, pem stats hsl, pem stats radius, pem stats subscriber, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2018-11-02 pem stats action(1)

pem stats dtos

NAME

dtos - Displays and resets PEM dtos statistics.

MODULE

pem stats

SYNTAX

Display statistics for the dtos component within the pem stats module using the syntax in the following section.

DISPLAY

show dtos

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the dtos component to display PEM dtos statistics. The statistics details are described below:

TAC Database

Specifies the number of TAC database queries and successful lookups since the last reset of the counter.

Customdb

Specifies the number of custom TAC database queries and successful lookups since the last reset of the counter.

Default/Licensed

Specifies the number of default/licensed TAC database queries and successful lookups since the last reset of the counter.

TCP Fingerprint

Specifies the number of TCP fingerprint queries and successful OS identification since the last reset of the counter.

User-agent

Specifies the number of user-agent queries and successful OS identification since the last reset of the counter.

You can reset the PEM dtos statistics using reset-stats command.

EXAMPLES

show dtos

Displays the PEM dtos statistics.

reset-stats dtos

Resets the PEM dtos statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gx, pem stats gy, pem stats radius, pem stats subscriber, pem stats hsl, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014. All rights reserved.

BIG-IP 2015-09-29 pem stats dtos(1)

pem stats gx

NAME

gx - Displays and resets PEM gx statistics.

MODULE

pem stats

SYNTAX

Display statistics for the gx component within the pem stats module using the syntax in the following section.

DISPLAY

show gx

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the gx component to display PEM gx statistics. The Current column represents count of Gx sessions that exist at the moment. The Max column represents maximal count of Gx sessions that ever existed at the same time. The Total column represents total count of Gx sessions that existed or messages that were sent or received since the start or the last statistics reset. The Retries column represents total count of retransmitted Gx messages. The Errors column represents total count of received Gx messages that were invalid or outgoing Gx messages that could not be sent. The statistics details are described below:

Gx sessions

Specifies the number of all Gx sessions established.

Provisioned Gx sessions

Specifies the number of provisioned Gx sessions.

Provisioning pending Gx sessions

Specifies the number sessions for which provisioning or creation over Gx has been initiated.

Not provisioned Gx sessions

Specifies the number of not provisioned Gx sessions, i.e. all existing Gx sessions that have not been provisioned.

Provisioning failed Gx sessions

Specifies the number of inactive Gx sessions for which provisioning or creation error happened.

Termination pending Gx sessions

Specifies the number of Gx sessions for which close is initiated.

Terminated Gx sessions

Specifies the total number of Gx sessions terminated since the last reset of the counter.

CCR Sent

Specifies the number of Credit-Control-Request (CCR) requests of all types sent.

CCR Initial Sent

Specifies the number of Credit-Control-Request (CCR) Initial requests sent since the last reset of the counter.

CCR Update Sent

Specifies the number of Credit-Control-Request (CCR) Update requests sent since the last reset of the counter.

CCR Usage Monitoring Sent

Specifies the number of Credit-Control-Request (CCR) requests with usage monitoring report sent.

CCR Application Start Sent

Specifies the number of Credit-Control-Request (CCR) requests with application start report sent.

CCR Application Stop Sent

Specifies the number of Credit-Control-Request (CCR) requests with application stop report sent.

CCR Termination Sent

Specifies the number of Credit-Control-Request (CCR) Termination requests sent since the last reset of the counter.

CCA Received

Specifies the number of Credit-Control-Answer (CCA) responses of all types received.

CCA Initial Received

Specifies the number of Credit-Control-Answer (CCA) Initial responses received since the last reset of the counter.

CCA Update Received

Specifies the number of Credit-Control-Answer (CCA) Update responses received since the last reset of the counter.

CCA Usage Monitoring Received

Specifies the number of Credit-Control-Answer (CCA) responses with usage monitoring report ack received.

CCA Termination Received

Specifies the number of Credit-Control-Answer (CCA) Termination responses received since the last reset of the counter.

RAR Received

Specifies the number of Re-Auth-Request (RAR) requests received.

RAA Sent

Specifies the number of Re-Auth-Answer (RAA) responses sent.

You can reset the PEM gx statistics using reset-stats command.

EXAMPLES

```
show gx
```

Displays the PEM gx statistics.

```
reset-stats gx
```

Resets the PEM gx statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gy, pem stats hsl, pem stats radius, pem stats subscriber, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

pem stats gy

NAME

gy - Displays and resets PEM gy statistics.

MODULE

pem stats

SYNTAX

Display statistics for the gy component within the pem stats module using the syntax in the following section.

DISPLAY

show gy

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the gy component to display PEM gy statistics. The Current column represents count of Gy sessions that exist at the moment. The Max column represents maximal count of Gy sessions that ever existed at the same time. The Total column represents total count of Gy that existed or messages that were sent or received since the start or the last statistics reset. The Retries column represents total count of retransmitted Gy messages. The Errors column represents total count of received Gy messages that were invalid or outgoing Gy messages that could not be sent. The statistics details are described below:

Gy sessions

Specifies the number of all Gy sessions established.

Late binding Gy sessions

Specifies the number of all late binding Gy sessions.

Provisioned Gy sessions

Specifies the number of provisioned Gy sessions.

Provisioning pending Gy sessions

Specifies the number sessions for which provisioning or creation over Gy has been initiated.

Not provisioned Gy sessions

Specifies the number of not provisioned Gy sessions, i.e. provisioning pending and provisioning failed Gy sessions.

Provisioning failed Gy sessions

Specifies the number of inactive Gy sessions for which provisioning or creation error happened.

Termination pending Gy sessions

Specifies the number of Gy sessions for which close is initiated.

Terminated Gy sessions

Specifies the total number of Gy sessions terminated since the last reset of the counter.

CCR Sent

Specifies the number of Credit-Control-Request (CCR) requests of all types sent.

CCR Initial Sent

Specifies the number of Credit-Control-Request (CCR) Initial requests sent since the last reset of the counter.

CCR Update Sent

Specifies the number of Credit-Control-Request (CCR) Update requests sent since the last reset of the counter.

CCR Termination Sent

Specifies the number of Credit-Control-Request (CCR) Termination requests sent since the last reset of the counter.

CCA Received

Specifies the number of Credit-Control-Answer (CCA) responses of all types received.

CCA Initial Received

Specifies the number of Credit-Control-Answer (CCA) Initial responses received since the last reset of the counter.

CCA Update Received

Specifies the number of Credit-Control-Answer (CCA) Update responses received since the last reset of the counter.

CCA Termination Received

Specifies the number of Credit-Control-Answer (CCA) Termination responses received since the last reset of the counter.

RAR Received
Specifies the number of Re-Auth-Request (RAR) requests received.

RAA Sent
Specifies the number of Re-Auth-Answer (RAA) responses sent.

You can reset the PEM gy statistics using `reset-stats` command.

EXAMPLES

`show gy`

Displays the PEM gy statistics.

`reset-stats gy`

Resets the PEM gy statistics.

OPTIONS

For information about the options that you can use with the command `show`, see `help show`.

SEE ALSO

`show`, `pem stats action`, `pem stats gx`, `pem stats hsl`, `pem stats radius`, `pem stats subscriber`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2017-05-15 pem stats gy(1)

pem stats hsl

NAME

`hsl` - Displays and resets PEM hsl statistics.

MODULE

`pem stats`

SYNTAX

Display statistics for the `hsl` component within the `pem stats` module using the syntax in the following section.

DISPLAY

`show hsl`

option:

(`default` | `exa` | `gig` | `kil` | `meg` | `peta` | `raw` | `tera` | `yotta` | `zetta`)

DESCRIPTION

You can use the `hsl` component to display PEM hsl statistics. The statistics details are described below:

Session Records

Specifies the number of Session-based records sent to each HSL endpoint since the last reset of the counter.

Flow Start Records

Specifies the number of Flow Start records sent to each HSL endpoint since the last reset of the counter.

Flow Interim Records

Specifies the number of Flow Interim records sent to each HSL endpoint since the last reset of the counter.

Flow Stop Records

Specifies the number of Flow Stop records sent to each HSL endpoint since the last reset of the counter.

Transaction Records

Specifies the number of HTTP Transaction records sent to each HSL endpoint since the last reset of the counter.

Records Skipped

Specifies the number of reporting records skipped from being sent the HSL endpoint since the last reset of the counter.

You can reset the PEM hsl statistics using `reset-stats` command.

EXAMPLES

`show hsl`

Displays the PEM hsl statistics.

reset-stats hsl

Resets the PEM hsl statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gx, pem stats gy, pem stats radius, pem stats subscriber, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2018-11-02 pem stats hsl(1)

pem stats hudnode-opt

NAME

hudnode-opt - Resets PEM hudnode-opt statistics.

MODULE

pem stats

SYNTAX

Reset statistics for the hudnode-opt component within the pem stats module using the syntax in the following section.

DESCRIPTION

You can use the hudnode-opt component to reset PEM hudnode-opt statistics.

EXAMPLES

reset-stats hudnode-opt

Resets the PEM hudnode-opt statistics.

OPTIONS

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, pem stats action, tmsb, pem stats gx

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2016. All rights reserved.

BIG-IP 2016-06-23 pem stats hudnode-opt(1)

pem stats multiple-ip

NAME

multiple-ip - Displays and resets PEM multiple IP statistics.

MODULE

pem stats

SYNTAX

Display statistics for the multiple-ip component within the pem stats module using the syntax in the following section.

DISPLAY

show multiple-ip

option:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the multiple-ip component to display PEM multiple IP statistics. The Total column represents total count of packets received. The Processed column represents total count of packets processed, i.e. packets that were successfully parsed, validated and initiated session creation, deletion or update. The Processed column is incremented regardless of success or failure of session creation, deletion or update and is incremented for retransmitted packets. The Retransmitted column represents total count of packets that were discarded, because they were retransmissions of already processed packets. The Error column represents total count of packets that could not be parsed or validated. The statistics details are described below:

Forwarded Flows

Specifies the number of flows that have been forwarded.

Sessions Merged

Specifies the number of sessions have been merged due to multiple IP.

You can reset the PEM multiple IP statistics using reset-stats command.

EXAMPLES

```
show multiple-ip
```

Displays the PEM multiple IP statistics.

```
reset-stats multiple-ip
```

Resets the PEM multiple IP statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gx, pem stats gy, pem stats hsl, pem stats subscriber, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2016. All rights reserved.

BIG-IP 2016-08-25 pem stats multiple-ip(1)

pem stats persistence

NAME

persistence - Displays and resets PEM persistence statistics.

MODULE

pem stats

SYNTAX

Display statistics for the persistence component within the pem stats module using the syntax in the following section.

DISPLAY

```
show persistence
```

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the persistence component to display PEM persistence statistics. The statistics details are described below:

Source IP

Number of times CARP persistence is based on source IP address.

Destination IP

Number of times CARP persistence is based on destination IP address.

HASH URI

Number of times CARP persistence is based on URI of the HTTP transaction.

HASH TCL

Number of times CARP persistence is based on TCL script.

Fallback Source IP

Number of times CARP persistence is based on fallback source IP address.

Fallback Destination IP

Number of times CARP persistence is based on fallback destination IP address.

You can reset the PEM persistence statistics using reset-stats command.

EXAMPLES

show persistence

Displays the PEM persistence statistics.

reset-stats persistence

Resets the PEM persistence statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gx, pem stats gy, pem stats radius, pem stats subscriber, pem stats hsl, pem stats dtos, pem stats tethering, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2015-06-05 pem stats persistence(1)

pem stats radius

NAME

radius - Displays and resets PEM radius statistics.

MODULE

pem stats

SYNTAX

Display statistics for the radius component within the pem stats module using the syntax in the following section.

DISPLAY

show radius

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the radius component to display PEM radius statistics. The Total column represents total count of packets received. The Processed column represents total count of packets processed, i.e. packets that were successfully parsed, validated and initiated session creation, deletion or update. The Processed column is incremented regardless of success or failure of session creation, deletion or update and is incremented for retransmitted packets. The Retransmitted column represents total count of packets that were discarded, because they were retransmissions of already processed packets. The Error column represents total count of packets that could not be parsed or validated. The statistics details are described below:

Accounting-Request START

Specifies the number of Accounting-Start packets.

Accounting-Request STOP

Specifies the number of Accounting-Stop packets.

Accounting-Request INTERIM_UPDATE

Specifies the number of Accounting-Interim packets.

Accounting-Request ACCOUNTING_ON

Specifies the number of Accounting-On packets.

Accounting-Request ACCOUNTING_OFF

Specifies the number of Accounting-Off packets.

You can reset the PEM radius statistics using reset-stats command.

EXAMPLES

show radius

Displays the PEM radius statistics.

reset-stats radius

Resets the PEM radius statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gx, pem stats gy, pem stats hsl, pem stats subscriber, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2015-12-08 pem stats radius(1)

pem stats sd

NAME

sd - Displays and resets PEM sd statistics.

MODULE

pem stats

SYNTAX

Display statistics for the sd component within the pem stats module using the syntax in the following section.

DISPLAY

show sd

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the sd component to display PEM sd statistics. The Current column represents count of Sd sessions that exist at the moment. The Max column represents maximal count of Sd sessions that ever existed at the same time. The Total column represents total count of Sd sessions that existed or messages that were sent or received since the start or the last statistics reset. The Retries column represents total count of retransmitted Sd messages. The Errors column represents total count of received Sd messages that were invalid or outgoing Sd messages that could not be sent. The statistics details are described below:

Sd sessions

Specifies the number of all Sd sessions established.

Provisioned Sd sessions

Specifies the number of provisioned Sd sessions.

Provisioning pending Sd sessions

Specifies the number sessions for which provisioning or creation over Sd has been initiated.

Not provisioned Sd sessions

Specifies the number of not provisioned Sd sessions, i.e. all existing Sd sessions that have not been provisioned.

Provisioning failed Sd sessions

Specifies the number of inactive Sd sessions for which provisioning or creation error happened.

Termination pending Sd sessions

Specifies the number of Sd sessions for which close is initiated.

Terminated Sd sessions

Specifies the total number of Sd sessions terminated since the last reset of the counter.

CCR Sent

Specifies the number of Credit-Control-Request (CCR) requests of all types sent.

CCR Initial Sent

Specifies the number of Credit-Control-Request (CCR) Initial requests sent since the last reset of the counter.

CCR Update Sent

Specifies the number of Credit-Control-Request (CCR) Update requests sent since the last reset of the counter.

CCR Usage Monitoring Sent

Specifies the number of Credit-Control-Request (CCR) requests with usage monitoring report sent.

CCR Application Start Sent

Specifies the number of Credit-Control-Request (CCR) requests with application start report sent.

CCR Application Stop Sent

Specifies the number of Credit-Control-Request (CCR) requests with application stop report sent.

CCR Termination Sent

Specifies the number of Credit-Control-Request (CCR) Termination requests sent since the last reset of the counter.

CCA Received

Specifies the number of Credit-Control-Answer (CCA) responses of all types received.

CCA Initial Received

Specifies the number of Credit-Control-Answer (CCA) Initial responses received since the last reset of the counter.

CCA Update Received

Specifies the number of Credit-Control-Answer (CCA) Update responses received since the last reset of the counter.

CCA Usage Monitoring Received

Specifies the number of Credit-Control-Answer (CCA) responses with usage monitoring report ack received.

CCA Termination Received

Specifies the number of Credit-Control-Answer (CCA) Termination responses received since the last reset of the counter.

RAR Received

Specifies the number of Re-Auth-Request (RAR) requests received.

RAA Sent

Specifies the number of Re-Auth-Answer (RAA) responses sent.

You can reset the PEM sd statistics using reset-stats command.

EXAMPLES

```
show sd
```

Displays the PEM sd statistics.

```
reset-stats sd
```

Resets the PEM sd statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gy, pem stats hsl, pem stats radius, pem stats subscriber, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2016-03-14 pem stats sd(1)

pem stats subscriber

NAME

subscriber - Displays and resets PEM subscriber statistics.

MODULE

pem stats

SYNTAX

Display statistics for the subscriber component within the pem stats module using the syntax in the following section.

DISPLAY

```
show subscriber
```

option:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the subscriber component to display PEM subscriber statistics. The Current column represents count of sessions or other objects that exist at the moment. The Max column represents maximal count of sessions or other objects that ever existed at the same time. The Total column represents total count of sessions or other

objects that existed since the start or the last statistics reset. Some entries support only the Total column. The statistics details are described below:

Active subscriber sessions

Specifies the number of all subscriber sessions. Note that you can use the db component in the sys module to configure tmm.pem.spm.maxsessionlimit to set the number of subscribers supported per processing unit (TMM). Then, the max number of the subscribers per device is set accordingly.

Subscriber sessions marked for deletion

Specifies the number of sessions that were marked for deletion, but not yet deleted.

Dynamic subscriber sessions without data from RADIUS

Specifies the number of sessions related to non-static subscribers that were neither created by RADIUS accounting start message nor updated by RADIUS accounting interim message. This counter is not incremented for a session related to static subscriber, even if this session is created dynamically.

Subscriber sessions by state/Waiting for provisioning

Specifies the number of sessions that are pending for provision.

Subscriber sessions by state/Waiting for provisioning Hold Timeout

Specifies the number of sessions that are within the provisioning hold time period.

Subscriber sessions by state/Provisioned

Specifies the number of provisioned sessions.

Subscriber sessions by state/Failed provisioning

Specifies the number of sessions, which provisioning failed. Usually this state is temporary and these sessions will try to be reprovisioned.

Subscriber sessions by origin/Sessions discovered via RADIUS

Specifies the number of sessions, which were created by RADIUS accounting start message.

Subscriber sessions by origin/Sessions discovered via DHCP

Specifies the number of sessions, which were created by DHCP message.

Subscriber sessions by origin/Sessions discovered from traffic

Specifies the number of sessions, which were created by client traffic that doesn't correspond to any existing session.

Subscriber sessions by origin/Sessions discovered via iRules

Specifies the number of sessions, which were created by iRules using PEM::session create or other command.

Subscriber sessions by origin/Pre-configured Static Subscriber Sessions

Specifies the number of sessions, which were created for static subscriber with pre-configured IP-addresses.

Subscriber sessions provisioning source/Sessions provisioned via Gx

Specifies the number of sessions, which were provisioned from PCRF via Gx interface.

Subscriber sessions provisioning source/Sessions provisioned via Sd

Specifies the number of sessions, which were provisioned from PCRF via Sd interface.

Subscriber sessions provisioning source/Sessions provisioned via iRules

Specifies the number of sessions, which were provisioned by iRule command.

Subscriber sessions provisioning source/Sessions provisioned from static subscriber database

Specifies the number of sessions which policies were defined on creation of the corresponding static subscriber.

Subscriber sessions provisioning source/Sessions not provisioned (pending and failed)

Specifies the number of not provisioned sessions.

Subscriber session update statistics/Session updates received

Specifies total count of session update attempts, e.g. by RADIUS accounting interim message, Gx Credit-Control-Request (CCR) Update request or iRule command.

Subscriber session update statistics/Session updates applied successfully

Specifies total count of session update attempts that were successful.

Subscriber session update statistics/Session updates received via RADIUS

Specifies total count of session update attempts by RADIUS accounting interim message.

Subscriber session update statistics/Session updates received via Gx

Specifies total count of session update attempts by Gx Credit-Control-Request updates (CCR-U) or Re-Auth-Requests updates (RAR).

Subscriber session update statistics/Session updates received via Sd

Specifies total count of session update attempts by Sd Credit-Control-Request updates (CCR-U) or Re-Auth-Requests updates (RAR).

Subscriber session update statistics/Session updates received via iRules

Specifies total count of session update attempts by iRule command.

Subscriber session deletion statistics/Total sessions fully deleted

Specifies total count of completely deleted sessions.

Subscriber session deletion statistics/Sessions replaced
Specifies total count of sessions that were deleted because they were replaced by other sessions.

Subscriber session deletion statistics/Provisioned sessions deleted
Specifies total count of deleted sessions that had been provisioned at the moment of deletion.

Subscriber session deletion statistics/Waiting for provisioning sessions deleted
Specifies total count of deleted sessions that were pending for provisioning at the moment of deletion.

Subscriber session deletion statistics/Failed provisioning sessions deleted
Specifies total count of deleted sessions that were in provisioning failed state at the moment of deletion.

Subscriber session deletion statistics/Sessions deleted via RADIUS
Specifies total count of sessions deleted by RADIUS accounting stop message.

Subscriber session deletion statistics/Sessions deleted via Gx
Specifies total count of sessions deleted by Gx RAR-t message.

Subscriber session deletion statistics/Sessions deleted due to Gx connectivity problems
Specifies total count of sessions deleted by the Fatal Grace Time timer associated with a diameter endpoint.

Subscriber session deletion statistics/Sessions deleted via Sd
Specifies total count of sessions deleted by Sd RAR-t message.

Subscriber session deletion statistics/Sessions deleted due to Sd connectivity problems
Specifies total count of sessions deleted by the Fatal Grace Time timer associated with a diameter endpoint.

Subscriber session deletion statistics/Sessions deleted via DHCP
Specifies total count of sessions deleted by DHCP release message.

Subscriber session deletion statistics/Sessions deleted due to Inactivity
Specifies total count of sessions deleted due to inactivity for time configured in the `tmm.pem.session.inactivitytimeout` variable.

Subscriber session deletion statistics/Sessions deleted via iRules
Specifies total count of sessions deleted by iRule command.

Subscriber session deletion statistics/Sessions deleted administratively
Specifies total count of sessions deleted interactively by operator.

Subscriber session deletion statistics/Sessions deleted due to static subscriber removal
Specifies total count of sessions deleted due to deletion of the corresponding static subscriber. This includes sessions that were created by RADIUS accounting start message, but deleted by deletion of static subscriber.

Subscriber session deletion statistics/Pre-configured static subscriber sessions deleted
Specifies total count of deleted sessions, which policies were defined on creation of the corresponding static subscriber. This includes sessions that were configured in static subscriber, but created and deleted by RADIUS accounting messages.

Multiple IP address subscriber sessions/Active sessions with single IP address
Specifies the number of sessions, which are associated with one IP address.

Multiple IP address subscriber sessions/Active sessions with multiple IP addresses
Specifies the number of sessions, which are associated with more than one IP address.

Multiple IP address subscriber sessions/Active sessions with both IPv4 and IPv6 addresses
Specifies the number of sessions, which are associated with one or multiple IPv4 addresses and one or multiple IPv6 addresses

Multiple IP address subscriber sessions/Active sessions with more than one IPv4 address
Specifies the number of sessions, which are associated with more than one IPv4 address.

Multiple IP address subscriber sessions/Active sessions with more than one IPv6 address
Specifies the number of sessions, which are associated with more than one IPv6 address.

Multiple IP address subscriber sessions/Total session IPv4 addresses
Specifies the number of IPv4 addresses associated with existing sessions. IP address uniqueness is not checked, so if the same IP address is added and removed several times, the total counter is incremented the each time address is added.

Multiple IP address subscriber sessions/Total session IPv6 addresses
Specifies the number of IPv6 addresses associated with existing sessions. IP address uniqueness is not checked, so if the same IP address is added and removed several times, the total counter is incremented the each time address is added.

Multiple IP address subscriber sessions/Sessions provisioned after hold timeout
Specifies the number of sessions, which were put on provision hold and were provisioned after the provision hold period finished.

Multiple IP address subscriber sessions/Session IP addresses added after session is provisioned
Specifies the total number of IP addresses that were added to already provisioned sessions. IP address uniqueness is not checked, so if the same IP address is added and removed several times, this counter is incremented the each time address is added.

Multiple IP address subscriber sessions/Session IP addresses removed without session delete
Specifies the total number of IP addresses that were removed from existing sessions without session deletion. This is possible for sessions associated with multiple IP addresses.

Exceptions, Errors, Other Statistics/Failed provisioning attempts
Specifies the aggregated number of failed provisioning attempts for all subscribers in the system since the last reset of the counter. For example, a provisioning attempt fails if a policy server (PCRF) returns an error, or does not respond for any reason.

Exceptions, Errors, Other Statistics/Subscriber Limit Exceeded
Specifies the count of the session creation failures, which are caused by exceeding the max number of subscribers supported by one processing unit (TMM).

Exceptions, Errors, Other Statistics/Session IP address limit exceeded
Specifies the count of the session creation or update failures, which are caused by exceeding the max number of IP addresses per session.

Exceptions, Errors, Other Statistics/Attempts to remove non-existent session via RADIUS
Specifies the count of attempts to remove session by RADIUS accounting stop message, which failed because the session was not found.

Exceptions, Errors, Other Statistics/Attempts to remove non-existent session via DHCP
Specifies the count of attempts to remove session by DHCP release message, which failed because the session was not found.

Exceptions, Errors, Other Statistics/Attempts to remove non-existent session via Gx
Specifies the count of attempts to remove session by Gx RAR-t message, which failed because the session was not found.

Exceptions, Errors, Other Statistics/Attempts to remove non-existent session via Sd
Specifies the count of attempts to remove session by Sd RAR-t message, which failed because the session was not found.

Exceptions, Errors, Other Statistics/Attempts to remove non-existent session via iRules
Specifies the count of attempts to remove session by iRule command, which failed because the session was not found.

Exceptions, Errors, Other Statistics/Subscriber sessions with activity log enabled Specifies the number of sessions for which the activity log was enabled.

You can reset the PEM subscriber statistics using reset-stats command.

EXAMPLES

```
show subscriber
```

Displays the PEM subscriber statistics.

```
reset-stats subscriber
```

Resets the PEM subscriber statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gx, pem stats gy, pem stats hsl, pem stats radius, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 pem stats subscriber(1)

pem stats tethering

NAME

tethering - Displays and resets PEM tethering statistics.

MODULE

pem stats

SYNTAX

Display statistics for the tethering component within the pem stats module using the syntax in the following section.

DISPLAY
show tethering
option:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the tethering component to display PEM tethering statistics. The statistics details are described below:

Subscribers Monitored

Specifies the number of subscribers being monitored for tethering since the last reset of the counter.

Subscribers Tethering

Specifies the number of subscribers tethering since the last reset of the counter.

Subscribers Tethering Maximum

Specifies the maximum number of subscribers that were found to be tethering at any point since the last reset of the counter.

You can reset the PEM tethering statistics using reset-stats command.

EXAMPLES

```
show tethering
```

Displays the PEM tethering statistics.

```
reset-stats tethering
```

Resets the PEM tethering statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, pem stats action, pem stats gx, pem stats gy, pem stats radius, pem stats subscriber, pem stats hsl, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014. All rights reserved.

BIG-IP 2014-10-27 pem stats tethering(1)

pem subscriber-attribute

NAME

subscriber-attribute - Configures subscriber attributes in Policy Enforcement Manager (PEM).

MODULE

pem

SYNTAX

Configure the subscriber-attribute component within the pem module using the syntax shown in the following sections.

CREATE/EDIT/MODIFY

```
create subscriber-attribute [name]
```

```
modify subscriber-attribute [name]
```

options:

```
app-service [[string] | none]
```

```
description [string]
```

```
export [disabled | enabled]
```

```
import [disabled | enabled]
```

```
well-known-attr-id [called-station-id | calling-station-id | imeisv | imsi | ipaddr | not-defined | subs-id | user-location-info | username]
```

```
edit subscriber-attribute [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list subscriber-attribute
```

```
list subscriber-attribute [ [ [name] | [glob] | [regex] ] ... ]
```

```
show running-config subscriber-attribute
```

```
show running-config subscriber-attribute [ [ [name] | [glob] | [regex] ] ... ]
```

options:

all-properties
non-default-properties
one-line
partition

DELETE

delete subscriber-attribute [name]

Note: You must remove all references to a subscriber-attribute before you can delete the subscriber-attribute.

DESCRIPTION

You can use the subscriber-attribute component to configure subscriber attribute definitions in Policy Enforcement Manager.

EXAMPLES

```
create subscriber-attribute 3gpp_imsi { import enabled export enabled well-known-attr-id imsi }
```

Creates a PEM subscriber attribute 3gpp_imsi with import enabled, export enabled, and well-known attribute id 'imsi'.

```
delete subscriber-attribute 3gpp_imsi
```

Deletes the PEM subscriber attribute named 3gpp_imsi.

```
list subscriber-attribute 3gpp_imsi
```

Displays the properties of the PEM subscriber attribute named 3gpp_imsi.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

export

Specifies whether the subscriber attribute can be exported (inserted) to the outgoing messages defined in pem protocol profile. This configuration is not applicable when referenced in pem protocol profile radius. It is applicable when referenced in pem protocol profile gx. The default value is enabled.

import

Specifies whether the subscriber attribute can be imported (parsed) from the incoming messages defined pem protocol profile. The default value is enabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the name of the administrative partition within which the subscriber attribute resides.

well-known-attr-id

Specifies an identifier of a well-known (build-in) subscriber attribute. The system provides a special handling for well-known subscriber attributes. For instance, most of the well-known attributes are included into session reporting records by default.

The options are:

called-station-id

The well-known subscriber attribute ID is called-station-id.

calling-station-id

The well-known subscriber attribute ID is calling-station-id.

imeisv

The well-known subscriber attribute ID is imeisv.

imsi The well-known subscriber attribute ID is imsi.

ipaddr

The well-known subscriber attribute ID is ipaddr. It can be ipv4 or ipv6 address.

not-defined

The well-known subscriber attribute ID is not defined. This is the default value.

subs-id

The well-known subscriber attribute ID is subs-id.

user-location-info

The well-known subscriber attribute ID is user-location-info.

username

The well-known subscriber attribute ID is username.

SEE ALSO

create, delete, edit, glob, list, modify, pem protocol profile gx, pem protocol profile radius, pem protocol

diameter-avp, pem protocol radius-avp, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 pem subscriber-attribute(1)

pem subscriber

NAME

subscriber - Configures subscribers for the Policy Enforcement Manager (PEM).

MODULE

pem

SYNTAX

Modify static subscriber component within PEM module using the syntax shown in the following sections.

CREATE/MODIFY

create subscriber [name]

modify subscriber [name]

options:

app-service [[string] | none]

ip-address-list [add | delete | replace-all-with] {
[ip address ...]

}
policies [add | delete | replace-all-with] {
[policy_name ...]

}
policies [default | none]

subscriber-id-type [dhcp | dhcp-custom | e164 | imsi | mac-address | mac-dhcp | nai | private]

edit subscriber [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list subscriber

list subscriber [[[name] | [glob] | [regex]] ...]

show running-config subscriber

show running-config subscriber [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

DELETE

delete subscriber [name]

Note: You must remove all references to a subscriber before you can delete the subscriber.

DESCRIPTION

You can use the subscriber component to configure subscriber definitions for the Policy Enforcement Manager.

Subscriber session IP addresses may optionally include a route domain ID. When specified, the route domain ID follows the IP address, after '%' separator. For example, 10.10.10.100%5 is an IP address in route domain 5.

A subscriber session may have multiple IP addresses in different route domains.

EXAMPLES

```
create subscriber 4085551212 { ip-address-list add { 10.10.10.2 10.10.10.3 } policies add { policy1 }  
subscriber-id-type imsi }
```

Creates a PEM subscriber 4085551212 with IP addresses 10.10.10.2 and 10.10.10.3, subscriber id type imsi, and a policy policy1.

```
create subscriber 4085551212 { ip-address-list add { 10.10.10.2%5 10.10.10.3%77 } policies add { policy1 }  
subscriber-id-type imsi }
```

Creates a PEM subscriber 4085551212 with IP addresses 10.10.10.2 and 10.10.10.3 in route domains 5 and 77 respectively, subscriber id type imsi, and a policy policy1.

```
create subscriber 4085551212 { ip-address-list add { 10.10.10.2%5 10.10.10.2%77 } policies add { policy1 }  
subscriber-id-type imsi }
```

Creates a PEM subscriber 4085551212 with two IP address entries that consist of the same IP address 10.10.10.2 and different route domain IDs: 5 and 77, subscriber id type imsi, and a policy policy1.

delete subscriber sub1

Deletes the subscriber named sub1.

list subscriber sub1

Displays the properties of the subscriber named sub1.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

name Specifies a unique subscriber identifier (up to 64 characters). This option is required for the commands create, delete, and modify.

ip-address-list

Adds, deletes, or replaces a list of the ip address to associate with the subscriber.

policies

Adds, deletes, or replaces a set of the policies to associate with the subscriber.

subscriber-id-type

Specifies the format to use for the subscriber id. The default value is imsi.

The options are:

dhcp For subscribers discovered via DHCP: an identifier comprises either Relay Option (option 82) for DHCPv4 based subscriber IDs or REMOTE-ID and SUBSCRIBER-ID Options (options 37 and 38) for DHCPv6, as configured in the corresponding DHCP profile.

dhcp-custom

For subscribers discovered via DHCP: an identifier created using a custom TCL snippet.

e164 A numbering plan that defines the format of an MSISDN international phone number (up to 15 digits). The number typically consists of three fields: country code, national destination code, and subscriber number.

imsi International Mobile Subscriber Identity. A globally unique code number that identifies a GSM, UMTS, or LTE mobile phone user.

mac-address

For subscribers discovered via DHCP: subscriber MAC address in a standard IEEE 802 format for MAC-48 (six groups of two hexadecimal digits, separated by colons ':').

mac-dhcp

For subscribers discovered via DHCP: a concatenation of mac-address and dhcp identifier as configured in the corresponding DHCP profile.

nai Network Access Identifier. A fully qualified network name in the form @; identifies a subscriber and the home network to which the subscriber belongs.

private

The subscriber id type is private for the given deployment.

SEE ALSO

create, delete, edit, glob, list, modify, pem forwarding-endpoint, pem interception-endpoint, pem listener, pem policy, pem profile diameter-endpoint, pem profile spm, pem reporting format-script, pem service-chain-endpoint, regex, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 pem subscriber(1)

pem subscribers

NAME

subscribers - Loads static subscribers for the Policy Enforcement Manager (PEM) from a file.

MODULE

pem

SYNTAX

Loads static subscribers from a file within the pem module using the syntax shown in the following sections.

LOAD

```
load subscribers file [filename]
```

DESCRIPTION

You can use the command `load pem subscribers` to load static subscribers definitions for the Policy Enforcement Manager (PEM). The maximum number of static subscribers allowed is $(2 * \text{sys db variable tmm.pem.spm.maxsessionlimit})$ or 100000, whichever is the lesser.

The static subscribers file is a csv file with the following fields: `„[,]*[,]+`. Each record can have zero IP address but must have at least one policy.

The maximum number of IP addresses per subscriber is set by sys db variable `tmm.pem.session.ip.addr.max`.

The maximum number of IPv4 addresses per subscriber is set by sys db variable `tmm.pem.session.ipv4.addr.max`.

The maximum number of IPv6 addresses per subscriber is set by sys db variable `tmm.pem.session.ipv6.addr.max`.

For example, these are the examples from such file:

```
subscriber1,e164,2,11.1.1.1,11.1.1.2,bronze,gold,silver
```

```
subscriber2,imsi,0,gold
```

The filename either absolute file name or just the base file name under folder: `/var/local/pem/subscribers/`

For more information about static subscriber, please refer to pem subscriber module.

EXAMPLES

```
load subscribers file my_ss_file
```

Loads static subscribers from file "my_ss_file" under the folder: `/var/local/pem/subscribers/`.

```
load subscribers file /shared/tmp/new_ss_file
```

Loads static subscribers from file "new_ss_file" under the folder: `/shared/tmp/`.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `modify`, `pem forwarding-endpoint`, `pem interception-endpoint`, `pem listener`, `pem policy`, `pem profile diameter-endpoint`, `pem profile spm`, `pem reporting format-script`, `pem service-chain-endpoint`, `regex`, `show`, `tmsb`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013. All rights reserved.

BIG-IP 2014-06-16 pem subscribers(1)

security

security analytics settings

NAME

settings - Configure global settings for security (AFM) analytics.

MODULE

security analytics

SYNTAX

Configure the settings component within the security analytics module using the syntax shown in the following sections.

MODIFY

```
modify settings
```

```
options:
```

```
acl-rules {  
  collect-client-ip [enabled | disabled]  
  collect-client-port [enabled | disabled]  
  collect-dest-ip [enabled | disabled]  
  collect-dest-port [enabled | disabled]  
  collect-server-side-stats [enabled | disabled]  
}
```

```

collected-stats-internal-logging [enabled | disabled]
collected-stats-external-logging [enabled | disabled]
dns {
  collect-client-ip [enabled | disabled]
}
dos-l2-l4 {
  collect-client-ip [enabled | disabled]
}
l3-l4-errors {
  collect-client-ip [enabled | disabled]
  collect-dest-ip [enabled | disabled]
}
publisher [name]
smtp-config [name]
stale-rules {
  collect [enabled | disabled]
}

```

DISPLAY
list settings

DESCRIPTION

Use the settings component to modify the settings for analytics entity collection for the AFM (advanced firewall) module.

EXAMPLES

```
modify settings acl-rules { collect-client-ip disabled }
```

Disables source/client IP analytics collection for ACL rules.

```
list settings
```

Displays analytics settings for AFM.

OPTIONS

acl-rules
Firewall (ACL) security statistics collection options.

collect-client-ip
Specifies whether source/client IP address should be collected for ACL rule matching.

collect-client-port
Specifies whether source/client port should be collected for ACL rule matching.

collect-dest-ip
Specifies whether the destination IP address should be collected for ACL rule matching.

collect-dest-port
Specifies whether the destination port should be collected for ACL rule matching.

collect-server-side-stats
Specifies whether server side statistics (source address translation information, self IP address and pool member address) should be collected for ACL rule matching.

collected-stats-internal-logging
Enables or disables the internal logging of the collected statistics.

collected-stats-external-logging
Enables or disables the external logging of the collected statistics.

dns DNS security statistics collection options.

collect-client-ip
Specifies whether source/client IP address should be collected for DNS security.

dos-l2-l4
Network DoS security statistics collection options.

collect-client-ip
Specifies whether source/client IP address should be collected for network layer's DoS security.

l3-l4-errors
Firewall errors statistics collection options.

collect-client-ip
Specifies whether source/client IP address should be collected for firewall errors.

collect-dest-ip
Specifies whether the destination IP address should be collected for firewall errors.

publisher
Specifies the external logging publisher used to send statistical data to one or more destinations.

smtp-config
Specifies the default SMTP configuration used for exporting CSV or PDF security analytics reports.

stale-rules

collect
Specifies whether statistics about all firewall rules should be collected in order to present information regarding rule staleness.

SEE ALSO

list, modify, show, tmsh, analytics network, analytics dos-l3, analytics dns-dos, analytics dns-protocol

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2013-10-15 security analytics settings(1)

security anti-fraud engine-update

NAME

engine-update - Runs or loads an Anti-fraud engine update.

MODULE

security anti-fraud

SYNTAX

Run or load the engine-update component within the security anti-fraud module using the syntax in the following sections:

LOAD

load engine-update

options:

file [filename]

If optional parameter file filename is specified in the load command, the command loads and installs engine update from local file instead of the cloud.

RUN

run engine-update

DISPLAY

list engine-update

options:

all-properties

current-version-create-datetime

download-available

install-datetime

install-user

message

non-default-properties

one-line

partition

progress-status

progress-status-datetime

last-update-check-datetime

readme

DESCRIPTION

You can use the engine-update component to run, load or display status of engine update.

EXAMPLES

list security anti-fraud engine-update

Displays the status of engine update.

OPTIONS

current-version-create-datetime

Displays the creation time of currently installed engine update version.

download-available

Displays whether new engine version is available for download from the cloud.

file Specifies the file name from which the engine update is going to be installed when using the load command. A full path should be specified.

install-datetime

Displays the time when engine update was installed.

install-user

Displays the name of the user who installed the last engine update.

message

Displays the message describing the failure status of engine update.

partition

Displays the administrative partition within which this object resides.

progress-status

Displays the engine update progress status.

progress-status-datetime

Displays the time when engine update progress status was last changed.

last-engine-check-datetime

Displays the time when last checked for engine update.

readme

Displays the Readme content for the current engine update.

SEE ALSO

list, security, security anti-fraud, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2017-01-17 security anti-fraud engine-update(1)

security anti-fraud profile

NAME

profile - Configures a Fraud Protection Service profile.

MODULE

security anti-fraud

SYNTAX

Configure the profile component within the security anti-fraud module using the syntax shown in the following sections.

CREATE/MODIFY

create profile [name]

modify profile [name]

options:

alert-client-side-caching [enabled | disabled]

alert-identifier [string]

alert-path [string]

alert-pool [[name] | none]

alert-publisher [[name] | none]

alert-token-header [string]

app-layer-encryption {

fail-open [enabled | disabled]

}

app-service [[string] | none]

auto-transactions {

bot-score [integer]

click-score [integer]

integrity-fail-score [integer]

min-mouse-move-count [integer]

min-mouse-over-count [integer]

min-report-score [integer]

min-time-to-request [integer]

not-human-score [integer]

strong-integrity {

hide-encrypted-parameters [enabled | disabled]

parameter [string]

}

tampered-cookie-score [integer]

time-fail-score [integer]

}

before-load-function [[string] | none]

blocking-page {

response-body [[string] | none]

response-headers [string]

}

[case-sensitive | case-insensitive]

```

cloud-service-pool [[name] | none]
config-location [string]
cookies {
  application [none | add | delete | replace-all-with] { [string] ... }
  base-domain {
apply [enabled | disabled]
exceptions [none | add | delete | replace-all-with] { [string] ... }
}
client-side [string]
client-side-lifetime [[integer] | session]
components-state [string]
components-state-lifetime [[integer] | session]
components-state-removal-protection [enabled | disabled]
encryption-disabled [string]
encryption-disabled-lifetime [[integer] | session]
encryption-disabled-removal-protection [enabled | disabled]
fingerprint [string]
fingerprint-lifetime [[integer] | session]
fingerprint-removal-protection [enabled | disabled]
html-field-obfuscation [string]
html-field-obfuscation-lifetime [[integer] | session]
malware-forensic [string]
malware-forensic-lifetime [[integer] | session]
malware-guid [string]
malware-guid-lifetime [[integer] | session]
malware-guid-removal-protection [enabled | disabled]
rules [string]
rules-lifetime [[integer] | session]
rules-removal-protection [enabled | disabled]
secure-alert [string]
secure-alert-lifetime [[integer] | session]
secure-alert-removal-protection [enabled | disabled]
secure-channel [string]
secure-channel-lifetime [[integer] | session]
secure-channel-removal-protection [enabled | disabled]
secure-mode [auto | disabled | enabled]
transaction-data [string]
transaction-data-lifetime [[integer] | session]
user-inspection [string]
user-name [string]
user-name-lifetime [[integer] | session]
user-name-removal-protection [enabled | disabled]
}
debug {
  console-log {
client-ips [none | add | delete | replace-all-with] { [string] ... }
user-agents [none | add | delete | replace-all-with] { [string] ... }
fingerprints [none | add | delete | replace-all-with] { [string] ... }
}
send-alert {
client-ips [none | add | delete | replace-all-with] { [string] ... }
user-agents [none | add | delete | replace-all-with] { [string] ... }
fingerprints [none | add | delete | replace-all-with] { [string] ... }
}
}
defaults-from [[name] | none]
description [[string] | none]
dummy-alert-html-maximum-length [integer]
encryption-staging-mode [enabled | disabled]
fingerprint {
  collect [enabled | disabled]
  location [string]
}
forensic {
  alert-path [string]
  client-domains [none | add | delete | replace-all-with] { [string] ... }
  cloud-config-path [string]
  cloud-forensics-mode [integer]
  cloud-remediation-mode [integer]
  continue-element [[string] | none]
  exe-location [string]
  html [[string] | none]
  self-post-location [string]
  skip-element [[string] | none]
  skip-path [string]
}
geolocation [enabled | disabled]
inject-main-javascript {
  [after | before]
  tag [string]
}
javascript-grace-threshold [integer]
javascript-location [string]
javascript-removal-location [string]
local-syslog-publisher [[name] | none]
malware {
  allowed-domains [none | add | delete | replace-all-with] { [string] ... }

```

```

    bait-check-generic [enabled | disabled]
    bait-location [string]
    blacklist-words [none | add | delete | replace-all-with] { [string] ... }
    detected-malware [none | add | delete | modify | replace-all-with] {
name [string] {
    baits [none | add | delete | modify | replace-all-with] {
        name [string] {
            data-before [string]
            data-inject [string]
            trigger-url {
name [string]
position [ alone | any | last ]
        }
    }
}
    blacklist-functions [none | add | delete | replace-all-with] { [string] ... }
    blacklist-js-words [none | add | delete | replace-all-with] { [string] ... }
    blacklist-urls [none | add | delete | replace-all-with] { [string] ... }
    blacklist-words [none | add | delete | replace-all-with] { [string] ... }
    browser-cache {
        blacklist-urls [none | add | delete | modify | replace-all-with] { [string] ... }
        whitelist-urls [none | add | delete | modify | replace-all-with] { [string] ... }
    }
    domain-availability {
        blacklist-urls [none | add | delete | modify | replace-all-with] { [string] ... }
        whitelist-urls [none | add | delete | modify | replace-all-with] { [string] ... }
    }
    dom-signatures [none | add | delete | modify | replace-all-with] {
        name [string] {
            attribute-name [[string] | none]
            hash-id [string]
            html-tag [[string] | none]
            match-type [ contains | is ]
            search-for [string]
            search-in [ all | attribute | html | js-global-variable | text ]
        }
    }
    generic-whitelist-words [none | add | delete | replace-all-with] { [string] ... }
}
    }
    domain-availability-urls [[string] | none]
    external-sources-targets [none | add | delete | replace-all-with] { [string] ... }
    flash-cookie-content [[string] | none]
    flash-cookie-location [string]
    flash-cookies [enabled | disabled]
    generic-whitelist-words [none | add | delete | replace-all-with] { [string] ... }
    inline-scripts-whitelist-signatures [none | add | delete | replace-all-with] { [string] ... }
    removed-scripts {
blacklist-functions [none | add | delete | replace-all-with] { [string] ... }
whitelist-functions [none | add | delete | replace-all-with] { [string] ... }
    }
    same-domain-scripts-validation-header [string]
    self-bait-header [string]
    source-integrity-location [string]
    web-rootkit {
blacklist-functions [none | add | delete | replace-all-with] { [string] ... }
whitelist-functions [none | add | delete | replace-all-with] { [string] ... }
    }
}
    mobilesafe {
        alert-custom-config [[string] | none]
        alert-threshold [integer]
        app-integrity {
custom-config [[string] | none]
[enabled | disabled]
android {
    score [integer]
    signature [[string] | none]
}
ios {
    hashes [none | add | delete | modify | replace-all-with] {
        value [string] {
            version [[string] | none]
        }
    }
    score [integer]
}
}
    general-custom-config [[string] | none]
    malware {
android {
    custom-malware [none | add | delete | modify | replace-all-with] {
        name [string] {
            package [string]
            score [integer]
        }
    }
}
}

```

```

custom-whitelist [none | add | delete | modify | replace-all-with] {
  name [string] {
    package [string]
  }
}
check-custom [enabled | disabled]
check-generic [enabled | disabled]
custom-config [[string] | none]
[enabled | disabled]
ios {
  custom-malware [none | add | delete | modify | replace-all-with] {
    name [string] {
      path [string]
      score [integer]
    }
  }
  custom-whitelist [none | add | delete | modify | replace-all-with] {
    name [string] {
      path [string]
    }
  }
}
behaviour-analysis {
  run [enabled | disabled]
  score [integer]
}
}
mitm {
certificate-custom-config [[string] | none]
dns-custom-config [[string] | none]
domains [none | add | delete | modify | replace-all-with] {
  name [string] {
    dns {
      ip-ranges [none | add | delete | replace-all-with] {address | address-address ... }
      spoofing-score [integer]
    }
    certificate {
      forging-score [integer]
      hash [string]
    }
  }
}
[enabled | disabled]
}
os-security {
android {
  untrusted-apps-score [integer]
  versions [none | add | delete | modify | replace-all-with] {
    priority [integer] {
      from [string]
      score [integer]
      to [string]
    }
  }
}
}
custom-config [[string] | none]
[enabled | disabled]
ios {
  versions [none | add | delete | modify | replace-all-with] {
    priority [integer] {
      from [string]
      score [integer]
      to [string]
    }
  }
}
}
rooting-jailbreak {
custom-config [[string] | none]
[enabled | disabled]
jailbreak-score [integer]
rooting-score [integer]
}
}
phishing {
  alert-path [string]
  allowed-elements [none | add | delete | replace-all-with] { [string] ...}
  allowed-referrers [none | add | delete | replace-all-with] { [string] ...}
  application-css [enabled | disabled]
  application-css-locations [none | add | delete | replace-all-with] { [string] ...}
  css-attribute-name [string]
  css-location [string]
  expiration-checks [enabled | disabled]
  image-location [string]
  inject-css-element {
[after | before]

```

```

tag [string]
  }
  inject-css-link {
[after | before]
tag [string]
  }
  inject-inline-javascript {
[after | before]
tag [string]
  }
  protected-elements [none | add | delete | replace-all-with] { [string] ...}
  referrer-checks [enabled | disabled]
}
referrer-info-header [string]
risk-engine-path [string]
risk-engine-publisher [[name] | none]
rules [none | add | delete | modify | replace-all-with] {
  event [auto-transaction | client-network-connection | client-side-missing-components | encryption-failure |
  generic-malware | mandatory-words | phishing | phishing-user | rat-detection | referrer-checks |
  server-side-missing-components | source-integrity | web-injection] {
action [block-user | forensic | inspection | redirect | remediation | route | web-service]
duration [integer]
enforce-policy [enforce | time-limited | unlimited]
min-score [integer]
publisher [[name] | none]
payload [[string] | none]
pool [[name] | none]
url [[string] | none]
  }
  }
  suggested-username-header [string]
  trigger-irule [enabled | disabled]
  urls [none | add | delete | modify | replace-all-with] {
    name [string] {
app-layer-encryption {
  add-decoy-inputs [enabled | disabled]
  auto-complete-block [enabled | disabled]
  auto-complete-whitelist-functions [none | add | delete | replace-all-with] { [string] ...}
  custom-encryption-function [[string] | none]
  [enabled | disabled]
  fake-strokes [enabled | disabled]
  full-ajax-encryption [enabled | disabled]
  hide-password-revealer [enabled | disabled]
  html-field-obfuscation [enabled | disabled]
  real-time-encryption [enabled | disabled]
  remove-element-ids [enabled | disabled]
  remove-event-listeners [enabled | disabled]
  stolen-creds [enabled | disabled]
  substitute-value-function [[string] | none]
}
auto-transactions {
  attach-ajax-payload-to-alerts [enabled | disabled]
  bot-score [integer]
  browser [enabled | disabled]
  click-score [integer]
  [enabled | disabled]
  full-ajax-integrity [enabled | disabled]
  integrity-fail-score [integer]
  integrity-fail-max-score [integer]
  min-mouse-move-count [integer]
  min-mouse-over-count [integer]
  min-report-score [integer]
  min-time-to-request [integer]
  non-browser [enabled | disabled]
  not-human-score [integer]
  strong-integrity [enabled | disabled]
  strong-integrity-user-functions [none | add | delete | replace-all-with] { [string] ...}
  submit-buttons [none | add | delete | replace-all-with] { [string] ...}
  tampered-cookie-score [integer]
  time-fail-score [integer]
}
before-load-function [[string] | none]
custom-alerts [none | add | delete | modify | replace-all-with] {
  name [string] {
    attach-request-part [enabled | disabled]
    component [auto-transactions | malware | mobilesafe | phishing]
    header-name [[string] | none]
    malware-name [[string] | none]
    message [[string] | none]
    search-in [client-ip | header | payload | query-string]
    value [[string] | none]
  }
}
description [string]
destination-urls [none | add | delete | replace-all-with] { [string] ...}
fallback-to-base-url [enabled | disabled]
include-query-string [enabled | disabled]

```

```

inject-javascript [enabled | disabled]
inject-javascript-removal {
  [after | before]
  tag [string]
}
inject-main-javascript {
  [after | before]
  tag [string]
}
login-response {
  status-code [[integer] | none]
  domain-cookie [[string] | none]
  exclude-string [[string] | none]
  header [[string] | none]
  include-string [[string] | none]
  validation [enabled | disabled]
}
malware {
  attach-html-to-alerts [enabled | disabled]
  auto-learn-form-tags [enabled | disabled]
  auto-learn-input-tags [enabled | disabled]
  auto-learn-script-tags [enabled | disabled]
  blocked-enter-key-detection [enabled | disabled]
  deferred-execution [enabled | disabled]
  domain-availability [enabled | disabled]
  enable-symbols [enabled | disabled]
  [enabled | disabled]
  external-injection [enabled | disabled]
  generic-malware [enabled | disabled]
  manual-count-form-tags [integer]
  manual-count-input-tags [integer]
  manual-count-script-tags [integer]
  password-exfiltration-detection [enabled | disabled]
  rat-detection [enabled | disabled]
  removed-scripts-detection [enabled | disabled]
  same-domain-scripts-validation [enabled | disabled]
  self-bait [enabled | disabled]
  source-integrity [enabled | disabled]
  vbklip-detection [enabled | disabled]
  visibility-check [enabled | disabled]
  visibility-check-items [none | add | delete | replace-all-with] { [string] ...}
  web-rootkit-detection [enabled | disabled]
  whitelist-dom-signatures [none | add | delete | replace-all-with] { [string] ...}
  whitelist-words [none | add | delete | replace-all-with] { [string] ...}
}
mobilesafe-encryption [enabled | disabled]
parameters [none | add | delete | modify | replace-all-with] {
  name [string] {
    ajax-mapping [string]
    attach-to-vtoken-report [enabled | disabled]
    check-integrity [enabled | disabled]
    encrypt [enabled | disabled]
    identify-as-username [enabled | disabled]
    method [GET | POST]
    mobilesafe-encrypt [enabled | disabled]
    mobilesafe-entangle [enabled | disabled]
    obfuscate [enabled | disabled]
    priority [integer]
    protect-by-selector [enabled | disabled]
    search-in [payload | query-string | any]
    substitute-value [enabled | disabled]
    type [explicit | wildcard]
  }
}
phishing {
  capture-users [enabled | disabled]
  copy-detection [enabled | disabled]
  css-protection [enabled | disabled]
  [enabled | disabled]
  field-types-to-send [none | add | delete | replace-all-with] { [string] ...}
  inject-css-element {
    [after | before]
    tag [string]
  }
  inject-css-link {
    [after | before]
    tag [string]
  }
  inject-inline-javascript {
    [after | before]
    tag [string]
  }
}
priority [integer]
type [explicit | wildcard]
}
}

```

```

users [add | delete | modify] {
  name [string] {
modes [add | delete] {
  mode [block | forensic | inspection | remediation] {
    duration [integer]
    enforce-policy [enforce | time-limited | unlimited]
    first-login-time [date]
  }
}
}
}
whitelist-custom-alerts [none | add | delete | replace-all-with] { [string] ...}

```

edit profile [[[name] | [glob] | [regex]] ...]

options:
 all-properties
 non-default-properties

DISPLAY

list profile

list profile [[[name] | [glob] | [regex]] ...]

show running-config profile

show running-config profile [[[name] | [glob] | [regex]] ...]

options:
 all-properties
 non-default-properties
 one-line
 partition
 recursive

DELETE

delete profile [name]

DESCRIPTION

You can use the profile component to create, modify, display, or delete an Anti-Fraud profile.

Note: The users property may be specified only for the commands modify, edit, and list and only when no other properties are specified. By default, users are not displayed.

Note: The first-login-time property of user modes may be specified only for the list command.

EXAMPLES

```
create profile my_antifraud_profile
```

Creates a custom Anti-Fraud profile named my_antifraud_profile with default parameters.

```
list profile
```

Displays the properties of all Anti-Fraud profiles.

OPTIONS

alert-client-side-caching

Specifies whether or not to cache the sent alerts in order to prevent multiple alerts from being sent to the dashboard.

alert-identifier

Specifies the ID of the customer in the dashboard.

alert-path

Specifies the BIG-IP URL path where the alert is sent. This path cannot be none and must start with '/'.

alert-pool

Specifies the name of the pool used when the system sends alerts.

alert-publisher

Specifies the name of the log publisher used for sending alerts originating from the BIG-IP. If only DPS is licensed, this publisher is used for reporting encryption failures.

alert-token-header

Specifies the name of the custom HTTP header in alerts for exchanging a random token between the client side and the BIG-IP.

app-layer-encryption

Specifies how the system performs Application layer encryption. With Application layer encryption, the system detects an attempt to steal and tamper with end-user passwords (or other protected information), and also prevents it by encrypting the protected information. You can configure the following options for Application layer encryption:

fail-open

Specifies, when enabled, that upon encryption error the system disables encryption in consecutive requests in the current session.

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

auto-transactions

Specifies how the system differentiates between human and automatic (bot) transactions. You can configure the following options for automatic transactions:

bot-score

Deprecated since v13.0.0. Please use bot-score in auto-transactions under urls instead. Specifies the score added to an alert that is triggered if the system determines that the client is a bot and not a human. The default is a score of 50.

click-score

Deprecated since v13.0.0. Please use click-score in auto-transactions under urls instead. Specifies the score added to an alert that is triggered if the min-mouse-over-count and min-mouse-move-count conditions are not met. The default is a score of 40.

integrity-fail-score

Deprecated since v13.0.0. Please use integrity-fail-score in auto-transactions under urls instead. Specifies the score added to an alert that is triggered if the system detects a difference between the actual parameter value and the expected value of a protected parameter sent after a user clicks a web form's Submit button. The default is a score of 40.

min-mouse-move-count

Deprecated since v13.0.0. Please use min-mouse-move-count in auto-transactions under urls instead. Specifies the minimum number of mouse movements necessary per page load in order for the system to consider the transaction to be of human origin. The default is 5 movements.

min-mouse-over-count

Deprecated since v13.0.0. Please use min-mouse-over-count in auto-transactions under urls instead. Specifies the minimum number of times the client's mouse is positioned over the Submit button in a web form in order for the system to consider the transaction to be of human origin. The default is 2 button interactions.

min-report-score

Deprecated since v13.0.0. Please use min-report-score in auto-transactions under urls instead. Specifies the lowest score necessary for the system to send an alert. The default value is 50.

min-time-to-request

Deprecated since v13.0.0. Please use min-time-to-request in auto-transactions under urls instead. Specifies the minimum amount of time (in seconds) permitted between when a web form is opened and the Submit button is clicked. The default is 2 seconds.

not-human-score

Deprecated since v13.0.0. Please use not-human-score in auto-transactions under urls instead. Specifies the score added to an alert that is triggered if the system only suspects that the client is a bot and not a human. The default is a score of 25.

strong-integrity

Specifies how the system performs strong integrity. You can configure the following options for strong integrity:

hide-encrypted-parameters

Deprecated since v14.1.0. Please use attach-to-vtoken-report under parameters instead. Specifies, when enabled, that JavaScript does not add the expected value of encrypted parameters to strong integrity parameter.

parameter

Deprecated since v14.1.0. Specifies the name of the HTTP parameter in POST requests added by JavaScript with the expected user-input data verified with physical input events.

tampered-cookie-score

Deprecated since v13.0.0. Please use tampered-cookie-score in auto-transactions under urls instead. Specifies the score added to an alert that is triggered if the system detects that the transaction-data cookie was tampered with. The default is a score of 50.

time-fail-score

Deprecated since v13.0.0. Please use time-fail-score in auto-transactions under urls instead. Specifies the score added to an alert that is triggered if the min-time-to-request condition is not met. The default is a score of 20.

before-load-function

Specifies the implementation of additional function to be run before JavaScript load, in the following format: function(configs){...}. Note: For certain advanced configurations, F5 support may provide relevant code to be entered here, please do not use it on your own.

blocking-page

Specifies information to display when the profile blocks a user account. You can configure the following options for blocking page:

response-body

Specifies the HTML code the system sends to the user whose account is blocked.

response-headers

Specifies the set of response headers that the system sends to the user whose account is blocked. Separate each header with a new line (Ctrl-V followed by Ctrl-J).

[case-sensitive | case-insensitive]

Specifies whether the profile treats protected URL paths as case sensitive, or not. The default value is case-insensitive. Note: If you create a profile, you can use either property, thereafter it becomes read

only. If the profile is case insensitive, the system stores protected URL paths in lowercase in the profile configuration.

`cloud-service-pool`
Specifies the name of the pool used by the system for various internal purposes, like signing Forensics tool.

`config-location`
Specifies the BIG-IP URL directory where the configuration for the injected JavaScript is located. The path here does not include the actual filename of the configuration for the injected JavaScript. This path cannot be none and must start with '/'.

`cookies`
Specifies names and lifetimes for the cookies that the system uses to optimize its detection of malware, data transactions, and phishing attacks on the web application. If you do not assign a name to a cookie, a random name is assigned. You can configure the following cookies:

`application`
Adds, deletes, or replaces a set of application cookies that will be removed if at least one of the protected cookies is missing.

`base-domain`
Specifies base domain settings for the cookies. You can configure the following options for base domain:

`apply`
Specifies, when enabled, that the system applies the cookies to the base domain.

`exceptions`
Adds, deletes, or replaces a set of exceptional base domains that take precedence when the system resolves the base domain from a host header.

`client-side`
Specifies the name of the cookie in which the system inserts plain text with a record about client side alerts already sent. This is done in order to prevent flooding the system with additional alerts if the page reloads.

`client-side-lifetime`
Specifies whether the client-side cookie is persistent, and if so, after how many minutes it expires.

`components-state`
Specifies the name of the cookie that verifies that the system's expected JavaScript can run successfully, and whether the system successfully decrypted configuration data arriving from server.

`components-state-lifetime`
Specifies whether the components-state cookie is persistent, and if so, after how many minutes it expires.

`components-state-removal-protection`
Enables or disables removal detection for the secure-alert cookie.

`encryption-disabled`
Specifies the name of the cookie that the system adds if the system fails to decrypt a password (to restore the original password as the user typed it), and the system forwards a request to the server and waits for a login failure response. In this case, the cookie does not encrypt the password on the next login attempt. This is used in situations where Application layer encryption is not possible (for example, if the user is using an old browser that cannot encrypt passwords).

`encryption-disabled-lifetime`
Specifies whether the encryption-disabled cookie is persistent, and if so, after how many minutes it expires.

`encryption-disabled-removal-protection`
Enables or disables removal detection for the encryption-disabled cookie.

`fingerprint`
Specifies the name of the cookie that contains fingerprint data.

`fingerprint-lifetime`
Specifies whether the fingerprint cookie is persistent, and if so, after how many minutes it expires.

`fingerprint-removal-protection`
Enables or disables removal detection for the fingerprint cookie.

`html-field-obfuscation`
Specifies the name of the cookie that the system sets to identify the fields that were created by HTML field obfuscation, in order to remove them from the request before sending it back to the web application, and to know which field names to decrypt.

`html-field-obfuscation-lifetime`
Specifies whether the html-field-obfuscation cookie is persistent, and if so, after how many minutes it expires.

`malware-forensic`
Specifies the name of the cookie that stores the essential response header values from the web

application to be sent to the user after he finishes or skips downloading and running Forensics tool on his host.

malware-forensic-lifetime

Specifies whether the malware-forensic cookie is persistent, and if so, after how many minutes it expires.

malware-guid

Specifies the name of the cookie set by JavaScript to a random string (12 chars long, not encrypted). The system sends this cookie value in a special alert to the dashboard in order to associate it with the logged in user.

malware-guid-lifetime

Specifies whether the malware-guid cookie is persistent, and if so, after how many minutes it expires.

malware-guid-removal-protection

Enables or disables removal detection for the malware-guid cookie.

rules

Specifies the name of the cookie that the system sets in order to perform the actions block-user, forensic, inspection, remediation, or redirect.

rules-lifetime

Specifies whether the rules cookie is persistent, and if so, after how many minutes it expires.

rules-removal-protection

Enables or disables removal detection for the rules cookie.

secure-alert

Specifies the name of the cookie that secures arrival of alerts originating from JavaScript to the dashboard.

secure-alert-lifetime

Specifies whether the secure-alert cookie is persistent, and if so, after how many minutes it expires.

secure-alert-removal-protection

Enables or disables removal detection for the secure-alert cookie.

secure-channel

Specifies the name of the cookie that the system sets when the system provides JavaScript with a public key for encryption operations. This cookie is used for the system to correlate incoming encrypted data with the private key when a request comes from the client.

secure-channel-lifetime

Specifies whether the secure-channel cookie is persistent, and if so, after how many minutes it expires.

secure-channel-removal-protection

Enables or disables removal detection for the secure-channel cookie.

secure-mode

Specifies the status of secure mode, to set 'Secure' flag or not for all FPS cookies.

auto Specifies that secure mode for FPS cookies will be set automatically depending on connection type. enabled for HTTPS (SSL) connections and disabled for HTTP connections. This is the default value.

disabled

Specifies that secure mode for FPS cookies will be disabled and FPS cookies will not have 'Secure' flag.

enabled

Specifies that secure mode for FPS cookies will be enabled and all FPS cookies will have 'Secure' flag.

transaction-data

Specifies the name of the cookie that contains information (such as mouse movement, clicks, and events) in encrypted format and sends that information to the system.

transaction-data-lifetime

Specifies whether the transaction-data cookie is persistent, and if so, after how many minutes it expires.

user-inspection

Specifies the name of cookie that is set once a user is identified in a web form submitted by the client and this user is enforced in inspection mode.

user-name

Specifies the name of the cookie with the username value after a username is identified in a request. This ensures that further transactions from the client are still associated with that user even if they do not include the username field.

user-name-lifetime

Specifies whether the user-name cookie is persistent, and if so, after how many minutes it expires.

user-name-removal-protection
Enables or disables removal detection for the user-name cookie.

debug
Specifies troubleshooting settings to add and filter debug logs of the system. Note: Only F5 support should configure this section, please do not use it on your own. F5 support can configure the following debug options:

console-log
Specifies when the system add prints to browser console. TMM logs are also enabled in such cases. F5 support can configure the following options for console log:

client-ips
Adds, deletes, or replaces a set of client IP addresses for which the system adds prints to browser console.

user-agents
Adds, deletes, or replaces a set of strings contained in user-agent header for which the system adds prints to browser console.

fingerprints
Adds, deletes, or replaces a set of strings contained in fingerprint data for which the system adds prints to browser console.

send-alert
Specifies when the system sends debug alerts to the dashboard. TMM logs are also enabled in such cases. F5 support can configure the following options for sending alerts:

client-ips
Adds, deletes, or replaces a set of client IP addresses for which the system sends debug alerts to the dashboard.

user-agents
Adds, deletes, or replaces a set of strings contained in user-agent header for which the system sends debug alerts to the dashboard.

fingerprints
Adds, deletes, or replaces a set of strings contained in fingerprint data for which the system sends debug alerts to the dashboard.

defaults-from
Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified.

description
User defined description.

dummy-alert-html-maximum-length
Specifies the maximum length of HTML attached to dummy alert.

encryption-staging-mode
Specifies, when enabled, that the system activates Anti-fraud encryption staging mode. If decrypted data differs from original data, an alert will be sent and original data will be used.

fingerprint
Specifies how the system collects fingerprint data. You can configure the following fingerprint options:

collect
Specifies, when enabled, that the system collects fingerprint data.

location
Specifies the BIG-IP URL location of the fingerprint JavaScript. This path cannot be none and must start with '/'.

forensic
Specifies how the system enforces scanning client host for malware (Forensics) and its removal (remediation). You can configure the following options for Forensics and remediation:

alert-path
Specifies the BIG-IP URL path for alerts from Forensics tool. This path cannot be none and must start with '/'.

client-domains
Adds, deletes, or replaces a set of client domains to be resolved by Forensics tool.

cloud-config-path
Specifies the BIG-IP URL path for requests from Forensics tool to cloud-service-pool. This path cannot be none and must start with '/'.

cloud-forensics-mode
Specifies the numeric value sent to cloud-service-pool to download Forensics tool.

cloud-remediation-mode
Specifies the numeric value sent to cloud-service-pool to download Forensics tool in remediation mode.

continue-element

Specifies the HTML element with continue option that replaces %SKIP_PART% in the entire html, when enforce-policy is enforce. Note: This property may be modified only when the DB variable antifraud.forensic.showgui has value enable.

exe-location

Specifies the BIG-IP URL path to download Forensics tool that also replaces %EXE_LOCATION% in the entire html. This path cannot be none and must start with '/'.

html Specifies the HTML code the system sends to the user after successful login with option to download Forensics tool. Note: This property may be modified only when the DB variable antifraud.forensic.showgui has value enable.

self-post-location

Specifies the BIG-IP URL path for self POST page opened by Forensics tool during scanning. This path cannot be none and must start with '/'.

skip-element

Specifies the HTML element with skip option that replaces %SKIP_PART% in the entire html, when enforce-policy is not enforce. Note: This property may be modified only when the DB variable antifraud.forensic.showgui has value enable.

skip-path

Specifies the BIG-IP URL path for skip / continue option that also replaces %SKIP_PATH% in both continue-element and skip-element (before their replacement in the entire html). This path cannot be none and must start with '/'.

geolocation

Specifies, when enabled, that the client collects geolocation data which will be sent as part of the alert data.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

inject-main-javascript

Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies where the system injects the main JavaScript. You can configure the following options for main JavaScript injection position:

[after | before]

Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies whether the system injects the main JavaScript after an opening tag or before a closing tag.

tag Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies the HTML tag for injection of the main JavaScript. This tag cannot be none.

javascript-grace-threshold

Specifies the maximum amount of time (in seconds) permitted between when a protected web page is loaded and its injected JavaScript activates.

javascript-location

Specifies the BIG-IP URL directory where the injected JavaScript is located. The path here does not include the actual filename of the injected JavaScript. This path cannot be none and must start with '/'.

javascript-removal-location

Specifies the BIG-IP URL location of the JavaScript removal detection location. This path cannot be none and must start with '/'.

local-syslog-publisher

DPS mode only. Specifies the name of the log publisher used for reporting encryption failures.

malware

Specifies how the system detects a malware attack on the web application. You can configure the following options for Malware protection:

allowed-domains

Adds, deletes, or replaces a set of whitelisted domains. The system does not send alerts on requests for URLs from these domains, even if the system detects malware injection on these domains.

bait-check-generic

Specifies, when enabled, that the system checks predefined baits. Note: The configured baits are checked anyway.

bait-location

Specifies the BIG-IP URL location of a file that acts as bait for attackers. This path cannot be none and must start with '/'.

blacklist-words

Deprecated since v13.0.0. Please use blacklist-js-words and blacklist-words in detected-malware instead. Adds, deletes, or replaces a set of words that are blacklisted if they appear in the web application's HTML or JavaScript code. If the system detects these words, the system generates a malware alert.

detected-malware

Adds, deletes, or replaces a set of malware detected by the system. You can configure the following options for each malware:

baits

Adds, deletes, or replaces a set of baits for this malware. You can configure the following options for each bait:

data-before

Specifies the HTML code that the malware searches and injects data-inject after it.

data-inject

Specifies the malicious code that the malware injects after data-before.

trigger-url

Specifies trigger URL settings for this bait. You can configure the following options for trigger URL:

name Specifies the URL pattern that triggers the malware to inject malicious code.

position

Specifies the position of this URL pattern in the query string of a bait request.

alone

Specifies that this trigger URL must be alone in the query string of a bait request.

any Specifies that the this trigger URL can be anywhere in the query string of a bait request. This is the default value.

last Specifies that the this trigger URL must be last in the query string of a bait request.

blacklist-functions

Adds, deletes, or replaces a set of regular expression patterns to detect functions that this malware can use when executing AJAX requests.

blacklist-js-words

Adds, deletes, or replaces a set of words that are blacklisted if they appear in the JavaScript code. If the system detects these words, the system generates a malware alert.

blacklist-urls

Adds, deletes, or replaces a set of regular expression patterns to detect URLs that this malware can use for AJAX requests and external scripts.

blacklist-words

Adds, deletes, or replaces a set of words that are blacklisted if they appear in the web application's HTML code. If the system detects these words, the system generates a malware alert.

browser-cache

Specifies how the system checks client network connection as targeted method. You can configure the following options for Browser cache:

blacklist-urls

Adds, deletes, or replaces a set of resources that are loaded by the malware.

whitelist-urls

Adds, deletes, or replaces a set of non-existent resources.

domain-availability

Specifies how the system checks client network connection as generic method. You can configure the following options for Domain availability:

blacklist-urls

Adds, deletes, or replaces a set of URLs that are not blocked by the malware.

whitelist-urls

Adds, deletes, or replaces a set of URLs that are blocked by the malware.

dom-signatures

Adds, deletes, or replaces a set of DOM signatures for this malware. You can configure the following options for each DOM signature:

attribute-name

Specifies the name of the attribute in which the pattern should be search for. Used only if search-in is attribute.

hash-id

Specifies unique ID that identifies this DOM signature in profile.

html-tag

Specifies the name of the HTML tag in which the pattern should be search for.

match-type

Specifies the type of DOM signature pattern matching.

contains

Specifies that this DOM signature pattern should be matched as partial match (not applicable when search-in is js-global-variable).

is Specifies that this DOM signature pattern should be matched as exact match.

`search-for`
Specifies the DOM signature pattern to search for.

`search-in`
Specifies search location for DOM signature.

`all` Specifies that this DOM signature should be searched in all locations.

`attribute`
Specifies that this DOM signature pattern should be searched only in an attribute with name `attribute-name`.

`html` Specifies that this DOM signature pattern should be searched only in HTML.

`js-global-variable`
Specifies that this DOM signature pattern should be searched only in JavaScript global variables (`match-type` contains not applicable in such case).

`text` Specifies that this DOM signature pattern should be searched only in text.

`generic-whitelist-words`
Deprecated since v15.0.0. Please use `whitelist-dom-signatures` in urls instead. Adds, deletes, or replaces a set of generic blacklisted words that are ignored.

`domain-availability-urls`
Deprecated since v13.0.0. Please use `blacklist-urls` and `whitelist-urls` in `domain-availability` under `detected-malware` instead. Specifies a JSON object containing URLs for which client network connectivity should be checked.

`external-sources-targets`
Adds, deletes, or replaces a set of HTML element types and their attributes for which external injections should be checked.

`flash-cookie-content`
Specifies the flash file (in hexadecimal format) used to allow JavaScript to access the Flash object on the client side. The default content is none. The length is limited to 64k.

`flash-cookie-location`
Specifies the BIG-IP URL location of the SWF file that JavaScript requests to get the Flash file. This path cannot be none and must start with '/'.

`flash-cookies`
Specifies, when enabled, that the system may use a Flash shared object (FSO) as a place to store an alternative malware cookie. This cookie tells the system, after a login attempt, that this user has malware, and the system sends an alert.

`generic-whitelist-words`
Deprecated since v13.0.0. Please use `generic-whitelist-words` in `detected-malware` instead. Adds, deletes, or replaces a set of generic blacklisted words that are ignored.

`inline-scripts-whitelist-signatures`
Adds, deletes, or replaces a set of signatures for allowed inline scripts. In case a signature appears as part of JavaScript inline script, the system does not count this script in the source integrity feature.

`removed-scripts`
Specifies how the system detects self-removed malicious scripts. You can configure the following options for removed scripts detection:

`blacklist-functions`
Adds, deletes, or replaces a set of functions that are used for detecting self-removed malicious scripts.

`whitelist-functions`
Adds, deletes, or replaces a set of functions that are NOT used for detecting self-removed malicious scripts.

`same-domain-scripts-validation-header`
Specifies the name of the custom HTTP header used to identify PING-PONG requests between JavaScript and BIG-IP for same-domain scripts validations. This name cannot be none.

`self-bait-header`
Specifies the name of the custom HTTP header used to identify self-bait requests from JavaScript to BIG-IP for malicious injections scan. This name cannot be none.

`source-integrity-location`
Specifies the BIG-IP URL path where the system collects information about the HTML source from multiple users. This path cannot be none and must start with '/'.

`web-rootkit`
Specifies how the system detects Web-RootKit malware. You can configure the following options for Web-RootKit detection:

`blacklist-functions`
Adds, deletes, or replaces a set of additional functions to be checked.

whitelist-functions

Adds, deletes, or replaces a set of native functions that are allowed to be overwritten.

mobilesafe

Specifies how the system detects and prevents phishing, Trojan, and pharming attacks on mobile devices in real time. You can configure the following options for mobile security:

alert-custom-config

Specifies alert custom configuration for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

alert-threshold

Specifies the minimal score for sending alerts from mobile devices.

app-integrity

Specifies how the system checks if the application on the mobile device has been tampered with. You can configure the following options for Application integrity:

custom-config

Specifies custom configuration of Application integrity for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

[enabled | disabled]

Enables or disables Application integrity.

android

Specifies Application integrity settings for Android platform. You can configure the following options for Android Application integrity:

score

Specifies Application integrity score for Android platform.

signature

Specifies signature of Android application (in hexadecimal format).

ios Specifies Application integrity settings for iOS platform. You can configure the following options for iOS Application integrity:

hashes

Adds, deletes, or replaces a set of iOS Application hashes (in base64-encoded format). You can configure the following options for iOS Application hash:

version

Specifies iOS Application version for this hash.

score

Specifies Application integrity score for iOS platform.

general-custom-config

Specifies general custom configuration for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

malware

Specifies how the system checks for malicious applications on the customer's mobile devices. You can configure the following options for Malware detection:

android

Specifies Malware detection settings for Android platform. You can configure the following options for Android Malware detection:

custom-malware

Adds, deletes, or replaces a custom set of checked malware for Android platform. You can configure the following options for each Android malware:

package

Specifies package of checked Android malware.

score

Specifies score for checked Android malware.

custom-whitelist

Adds, deletes, or replaces a custom set of whitelist applications for Android platform. You can configure the following options for each whitelist Android application:

package

Specifies package of whitelist Android application.

check-custom

Enables or disables custom malware check.

check-generic

Enables or disables generic malware check.

custom-config

Specifies custom configuration of Malware detection for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

[enabled | disabled]

Enables or disables Malware detection.

ios Specifies Malware detection settings for iOS platform. You can configure the following options for iOS Malware detection:

custom-malware

Adds, deletes, or replaces a custom set of checked malware for iOS platform. You can configure the following options for each iOS malware:

path Specifies path of checked iOS malware.

score

Specifies score for checked iOS malware.

custom-whitelist

Adds, deletes, or replaces a custom set of whitelist applications for iOS platform. You can configure the following options for each whitelist iOS application:

path Specifies path of whitelist iOS application.

behaviour-analysis

Specifies how the system checks for suspicious behavior and characteristics on all applications on the customer's mobile devices. You can configure the following options for behavior analysis:

run Enables or disables behaviour analysis run.

score

Specifies score for behavior analysis.

mitm Specifies how the system checks the defined domains for DNS Spoofing and Certificate Forging on customer devices. You can configure the following options for Man-in-the-middle detection:

certificate-custom-config

Specifies custom configuration of Certificate forging detection for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

dns-custom-config

Specifies custom configuration of DNS spoofing detection for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

domains

Adds, deletes, or replaces a set of domains for Man-in-the-middle detection. You can configure the following options for a MITM domain:

dns Specifies DNS spoofing detection settings for this domain. You can configure the following options for DNS spoofing detection:

ip-ranges

Adds, deletes, or replaces a set of IP address ranges for DNS spoofing detection.

spoofing-score

Specifies score for DNS spoofing detection.

certificate

Specifies Certificate forging detection settings for this domain. You can configure the following options for Certificate forging detection:

forging-score

Specifies score for Certificate forging detection.

hash Specifies certificate hash.

[enabled | disabled]

Enables or disables Man-in-the-middle detection.

os-security

Specifies how the system checks the customer's mobile devices for old, unsupported, and unpatched operation system (OS) versions. You can configure the following options for OS security:

android

Specifies OS security settings for Android platform. You can configure the following options for Android OS security:

versions

Adds, deletes, or replaces an ordered set of version ranges for Android platform. You can configure the following options for Android version range:

from Specifies Android version number from which OS is unpatched.

priority

Specifies a unique ordinal number for Android version range in the set. This option is required for the operations add, delete, modify, and replace-all-with.

score

Specifies score for Android version range.

to Specifies Android version number to which OS is unpatched.

custom-config

Specifies custom configuration of OS security for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

[enabled | disabled]

Enables or disables OS security.

ios Specifies OS security settings for iOS platform. You can configure the following options for iOS OS security:

versions

Adds, deletes, or replaces an ordered set of version ranges for iOS platform. You can configure the following options for iOS version range:

from Specifies iOS version number from which OS is unpatched.

priority

Specifies a unique ordinal number for iOS version range in the set. This option is required for the operations add, delete, modify, and replace-all-with.

score

Specifies score for iOS version range.

to Specifies iOS version number to which OS is unpatched.

untrusted-apps-score

Specifies score for untrusted applications.

rooting-jailbreak

Specifies how the system checks customer's mobile devices to determine if they are rooted / jailbroken. You can configure the following options for Rooting / Jailbreak detection:

custom-config

Specifies custom configuration of Rooting / Jailbreak detection for SDK forward compatibility. Note: For certain advanced configurations, F5 support may provide a relevant string to be entered here, please do not use it on your own.

[enabled | disabled]

Enables or disables Rooting / Jailbreak detection.

jailbreak-score

Specifies score for jailbreak on iOS platform.

rooting-score

Specifies score for rooting on Android platform.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

phishing

Specifies how the system detects a phishing attempt. You can configure the following options for phishing site detection:

alert-path

Specifies the BIG-IP URL path for alerts from the phishing inline script. This path cannot be none and must start with '/'.

allowed-elements

Adds, deletes, or replaces a set of URLs in requests for which the system does not verify (check) the referrer header value.

allowed-referrers

Adds, deletes, or replaces a set of domain names that are allowed to appear in the referrer header when requesting protected resources.

application-css

Specifies, when enabled, that the system injects the CSS content to the existing application CSS files.

application-css-locations

Adds, deletes, or replaces a set of server URL locations of the application CSS files, used when application-css is enabled.

css-attribute-name
Specifies the attribute name as part of the CSS content. This name cannot be none.

css-location
Specifies the BIG-IP URL location of the CSS file, used when application-css is disabled. Injecting JavaScript protects the web application against phishing attempts because even if an attacker removes the injected JavaScript from the copied web page, the CSS element is not modified, and this triggers an alert. This path cannot be none and must start with '/'.

expiration-checks
Specifies, when enabled, that the system sends an alert if expired JavaScript engine files are used, as this is an indication of a phishing attack.

image-location
Specifies the BIG-IP URL location of the 1x1 pixel image file. If an attacker copies a web page with this image, it most likely lacks the JavaScript, and this triggers an alert. This path cannot be none and must start with '/'.

inject-css-element
Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies where the system injects the CSS element. You can configure the following options for CSS element injection position:

[after | before]
Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies whether the system injects the CSS element after an opening tag or before a closing tag.

tag Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies the HTML tag for injection of the CSS element. This tag cannot be none.

inject-css-link
Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies where the system injects the CSS link, when application-css is disabled. You can configure the following options for CSS link injection position:

[after | before]
Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies whether the system injects the CSS link after an opening tag or before a closing tag.

tag Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies the HTML tag for injection of the CSS link. This tag cannot be none.

inject-inline-javascript
Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies where the system injects the phishing inline script and image. You can configure the following options for phishing inline script and image injection position:

[after | before]
Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies whether the system injects the phishing inline script and image after an opening tag or before a closing tag.

tag Deprecated since v12.1.3 (excluding v13.0.0). Please use same configuration in a specific URL instead. Specifies the HTML tag for injection of the phishing inline script and image. This tag cannot be none.

protected-elements
Adds, deletes, or replaces a set of URLs in requests for which the system verifies (checks) the referrer header value. You can use wildcards, for example *.gif.

referrer-checks
Specifies, when enabled, that the system verifies (checks) requests coming to the web application for resources from different domains.

referrer-info-header
Specifies the name of the custom HTTP header used by client side to communicate referrer and view identifier to BIG-IP.

risk-engine-path
Specifies the BIG-IP URL path to where a risk-engine report is sent by client. This path cannot be none and must start with '/'.

risk-engine-publisher
Specifies the name of the log publisher used for reports to a Risk engine.

rules
Adds, deletes, or replaces a set of rules used by the system to perform actions upon detected events. You can configure the following options for each rule:

action
Specifies the type of the action that the system performs when this event is detected. The options are:

block-user
Specifies that the system adds the user with block mode to be enforced from the next login.

forensic

Specifies that the system adds the user with forensic mode to be enforced from the next login.

inspection

Specifies that the system adds the user with inspection mode to be enforced from the next login.

redirect

Specifies that the system redirects the next request to a specific web page.

remediation

Specifies that the system adds the user with remediation mode to be enforced from the next login.

route

Deprecated in v13.1.0. Specifies that the system routes to a specific pool all subsequent requests for a specific time.

web-service

Specifies that the system sends a POST request to a specific Web service.

duration

Specifies number of minutes during which the system performs the action block-user, forensic, inspection, remediation or route.

enforce-policy

Specifies enforcement policy for the action block-user, forensic, inspection or remediation. The options are:

enforce

Specifies that the system adds the user mode with the enforce policy.

time-limited

Specifies that the system adds the user mode with the time-limited policy.

unlimited

Specifies that the system adds the user mode with the unlimited policy.

event

Specifies a unique event for the rule. This option is required for the operations create, delete, modify, and replace-all-with. The options are:

auto-transaction

Specifies that the action is performed when the system detects automatic (bot) transaction.

client-network-connection

Specifies that the action is performed when the system detects that client network connectivity is blocked.

client-side-missing-components

Specifies that the action is performed when the system detects missing components on the client side.

encryption-failure

Specifies that the action is performed when the system fails to decrypt a password.

generic-malware

Specifies that the action is performed when the system detects generic malware.

mandatory-words

Specifies that the action is performed when the system detects that mandatory words are changed in the page.

phishing

Specifies that the action is performed when the system detects a phishing attempt.

phishing-user

Specifies that the action is performed when the system detects a user attacked by a phishing attempt.

rat-detection

Specifies that the action is performed when the system detects a Remote Access Trojan (RAT) on a client web browser.

referrer-checks

Specifies that the action is performed when the system detects a request from a different domain by the referrer header.

server-side-missing-components

Specifies that the action is performed when the system detects missing components on the BIG-IP.

source-integrity

Specifies that the action is performed when the system detects a mismatch of the URL's HTML source code.

web-injection

Specifies that the action is performed when the system detects an attempt to inject malware.

min-score

Specifies the lowest score of this event necessary for the system to perform the action.

payload

Specifies the payload for the web-service action.

pool Specifies the name of the pool for the route action.

publisher

Specifies the name of the log publisher for the web-service action.

url Specifies the URL for the action redirect or web-service.

suggested-username-header

Specifies the name of the custom HTTP header in AJAX requests added by JavaScript with a username value identified on the client side.

trigger-irule

Specifies, when enabled, that the system activates Anti-fraud iRule events. The default value is disabled.

urls Adds, deletes, or replaces a set of URLs in the web application that are protected by the system. You can configure the following options for a protected URL:

app-layer-encryption

Specifies how the system performs Application layer encryption for this URL. With Application layer encryption, the system detects an attempt to steal and tamper with end-user passwords (or other protected information), and also prevents it by encrypting the protected information. You can configure the following options for Application layer encryption:

add-decoy-inputs

Specifies, when enabled, that the system randomly and continuously generates and removes decoy fields that are added to the web page, thus making it harder for an attacker to identify sensitive information with either JavaScript or a proxy. In order to enable it, you must first enable html-field-obfuscation.

auto-complete-block

Specifies, when enabled, that the system prevents auto-complete functionality in browser.

auto-complete-whitelist-functions

Specifies a list of customer-specific global functions that require access to the value of a parameter with substitute-value enabled.

custom-encryption-function

Specifies the name or implementation of custom encryption function to be run instead of built-in encryption.

[enabled | disabled]

Specifies whether the system protects this URL with Application layer encryption, and sends an alert if an attacker attempts to breach Application layer encryption for this URL, or not.

fake-strokes

Specifies, when enabled, that the system protects against in-browser key loggers by generating fake keyboard events.

full-ajax-encryption

Specifies, when enabled, that the system encrypts the full AJAX payload.

hide-password-revealer

Specifies, when enabled, that the system hides the password revealer icon found in web pages.

html-field-obfuscation

Specifies, when enabled, that the system encrypts the names of defined fields on the client, and then decrypts them back to the original names on the BIG-IP.

real-time-encryption

Specifies, when enabled, that the system encrypts passwords as they are typed (even before the user clicks the Submit button in a web form).

remove-element-ids

Specifies, when enabled, that the system removes the ID attribute from the fields in a web form. In order to enable it, you must first enable html-field-obfuscation.

remove-event-listeners

Specifies, when enabled, that the system removes event listeners from the encrypted fields in a web form.

stolen-creds

Specifies, when enabled, that the system examines whether the user was trying to use a fabricated password.

substitute-value-function

Specifies a JavaScript function that receives the real password as an argument and returns a

fake value.

auto-transactions

Specifies how the system protects this URL from automatic (bot) transactions. You can configure the following options for Automated transactions detection:

attach-ajax-payload-to-alerts

Specifies whether to attach the actual AJAX payload to alerts. Use the DB variable `antifraud.antifraud.maxalertrequestsize` to limit the attached payload size.

bot-score

Specifies the score added to an alert that is triggered if the system determines that the client is a bot and not a human. The default is a score of 50.

browser

Specifies, when enabled, that the system looks for bot automation performed within the browser.

click-score

Specifies the score added to an alert that is triggered if the `min-mouse-over-count` and `min-mouse-move-count` conditions are not met. The default is a score of 40.

[enabled | disabled]

Specifies whether the system protects this URL against non-human transactions, and sends an alert if the system detects a non-human transaction attempt for this URL, or not.

full-ajax-integrity

Specifies, when enabled, that the system verifies whether the full AJAX payload was changed by malware when it left the browser for the server.

integrity-fail-score

Specifies the score added to an alert that is triggered if the system detects a difference between the actual parameter value and the expected value of a protected parameter sent after a user clicks a web form's Submit button. The default is a score of 40.

integrity-fail-max-score

Specifies the maximal score added to an alert that is triggered if the system detects a difference between the actual parameter value and the expected value of a protected parameter sent after a user clicks a web form's Submit button. The default is a score of 100

min-mouse-move-count

Specifies the minimum number of mouse movements necessary per page load in order for the system to consider the transaction to be of human origin. The default is 5 movements.

min-mouse-over-count

Specifies the minimum number of times the client's mouse is positioned over the Submit button in a web form in order for the system to consider the transaction to be of human origin. The default is 2 button interactions.

min-report-score

Specifies the lowest score necessary for the system to send an alert. The default value is 50.

min-time-to-request

Specifies the minimum amount of time (in seconds) permitted between when a web form is opened and the Submit button is clicked. The default is 2 seconds.

non-browser

Specifies, when enabled, that the system looks for bot automation performed not within the browser.

not-human-score

Specifies the score added to an alert that is triggered if the system only suspects that the client is a bot and not a human. The default is a score of 25.

strong-integrity

Specifies, when enabled, that Enhanced Data Integrity is active. When Enhanced Data Integrity is active, the system detects a difference between the actual parameter value and the expected value of a protected parameter verified with physical input events.

strong-integrity-user-functions

Adds, deletes, or replaces a set of configures a list of customer functions that change a parameter value protected by Enhanced Data Integrity.

submit-buttons

Adds, deletes, or replaces a set of non-standard Submit buttons found in forms of the web application. You can specify the name, or the CSS syntax (ID, class, or tagname) for each button.

tampered-cookie-score

Specifies the score added to an alert that is triggered if the system detects that the `transaction-data` cookie was tampered with. The default is a score of 50.

time-fail-score

Specifies the score added to an alert that is triggered if the `min-time-to-request` condition is not met. The default is a score of 20.

custom-alerts

Adds, deletes, or replaces a set of user-defined alerts sent by the system upon searches in

different parts of the request. You can configure the following options for each user-defined alert:

attach-request-part

Specifies whether to attach the original client-side request to this alert.

component

Specifies the alert component that the system sends in this alert. Select either: malware (the default value), phishing, auto-transactions, or mobilesafe.

header-name

Specifies a header name in which the system searches for the value when search-in is header.

malware-name

Specifies the malware detected by this alert when component is malware.

message

Specifies the user-defined message that the system sends in this alert.

search-in

Specifies the part of the request where the system must find the value to send this alert.

Note: If you create a user-defined alert, you can use either request part, thereafter it becomes read only.

client-ip

Specifies that the systems sends this alert if the client IP address equals to the value.

header

Specifies that the systems sends this alert if the header-name header contains the value.

payload

Specifies that the systems sends this alert if the request payload contains the value.

query-string

Specifies that the systems sends this alert if the URL query string contains the value.

value

Specifies a value that the system searches for in the search-in part of the request. The default value is none, which means that the system searches for any value.

before-load-function

Specifies the implementation of additional function to be run before JavaScript load, in the following format: function(configs){...}. Note: For certain advanced configurations, F5 support may provide relevant code to be entered here, please do not use it on your own.

description

Specifies an optional description of this URL.

destination-urls

Specifies a list of destination URLs for requests from SPA URLs/Views.

fallback-to-base-url

Specifies if a request to a non-configured view should use same configuration as the base URL or disable FPS for that request.

include-query-string

Specifies, when enabled, that the system includes query string of URLs to match this wildcard expression. The default value is disabled.

inject-javascript

Enables or disables JavaScript injection into responses to this URL. The default value is enabled.

inject-main-javascript

Specifies where the system injects the main JavaScript. You can configure the following options for main JavaScript injection position:

[after | before]

Specifies whether the system injects the main JavaScript after an opening tag or before a closing tag.

tag Specifies the HTML tag for injection of the main JavaScript. This tag cannot be none.

inject-javascript-removal

Specifies where the system injects the JavaScript removal detection image. You can configure the following options for JavaScript removal detection image injection position:

[after | before]

Specifies whether the system injects the JavaScript removal detection image after an opening tag or before a closing tag.

tag Specifies the HTML tag for injection of the JavaScript removal detection image. This tag cannot be none.

login-response

Specifies validation criteria on the response of this URL when it is Login page. You must configure at least one of them. If you configure more than one validation criteria, then all the criteria must be fulfilled for successful login. You can configure the following Login page properties:

status-code

Specifies an HTTP response status code that the server must return to the user upon successful login.

domain-cookie

Specifies a defined domain cookie that the successful response to the login URL must include.

exclude-string

Specifies a string that should NOT appear in the successful response to the login URL.

header

Specifies a header name and value that the successful response to the login URL must match.

include-string

Specifies a string that should appear in the successful response to the login URL.

validation

Enables or disables successful login validation.

malware

Specifies when the system detects attempts of attackers to inject malware in the URL. You can configure the following options for Malware detection:

attach-html-to-alerts

Specifies, when enabled, that the system attaches forensics information along with the alerts.

auto-learn-form-tags

Specifies, when enabled, that the system learns the number of HTML form tags that appear in the URL. In order to enable it, you must first enable source-integrity.

auto-learn-input-tags

Specifies, when enabled, that the system learns the number of HTML input tags that appear in the URL. In order to enable it, you must first enable source-integrity.

auto-learn-script-tags

Specifies, when enabled, that the system learns the number of HTML script tags that appear in the URL. In order to enable it, you must first enable source-integrity.

blocked-enter-key-detection

Specifies, when enabled, that the system detects blocked "Enter" key.

deferred-execution

Specifies, when enabled, that the system detects deferred execution attack.

domain-availability

Specifies, when enabled, that the system checks that client network connectivity is not blocked by malware.

enable-symbols

Specifies, when enabled, that the system looks for malware strings (signatures) within JavaScript.

[enabled | disabled]

Specifies whether the system protects this URL against injected malware, and sends an alert if this URL is detected to have malware, or not.

external-injection

Specifies, when enabled, that the system detects malicious scripts injected from domains not in the profile's allowed-domains.

generic-malware

Specifies, when enabled, that the system applies the detection of generic malware, using honeypots.

manual-count-form-tags

Specifies the number of HTML forms that appear in the URL.

manual-count-input-tags

Specifies the number of HTML inputs that appear in the URL.

manual-count-script-tags

Specifies the number of HTML scripts that appear in the URL.

password-exfiltration-detection

When enabled, the system detects attempts to steal the user's password in the web browser. An alert is triggered if such an attempt is detected.

rat-detection

Specifies, when enabled, that the system checks for Remote Access Trojans (RATs) on clients' web browsers.

removed-scripts-detection

Specifies, when enabled, that the system detects malicious scripts that removed their own injection from the DOM.

same-domain-scripts-validation

Specifies, when enabled, that the system detects malicious responds to same-domain scripts.

self-bait

Specifies, when enabled, that the system scans the original source code of the page for malicious injections.

source-integrity

Specifies, when enabled, that the system verifies that the URL's HTML source code matches the HTML code sent from the server. The source integrity feature counts script tags that are external (with src) and inline (without src).

vbklip-detection

Specifies, when enabled, that the system checks for VBKlip malware.

visibility-check

Specifies, when enabled, that the system searches HTML pages for words from visibility-check-items.

visibility-check-items

Adds, deletes, or replaces a set of words that must appear in the web site's HTML pages and may not be changed. If these words are changed, the system sends an alert.

web-rootkit-detection

Specifies, when enabled, that the system detects malware that overwrites native browser functions.

whitelist-dom-signatures

Adds, deletes, or replaces a set of hash-IDs of DOM signatures that are permitted to appear in requests for this URL, even though they are otherwise blacklisted by the system for other URLs.

whitelist-words

Deprecated since v15.0.0. Please use 'whitelist-dom-signatures' configuration instead. Adds, deletes, or replaces a set of words that are permitted to appear in requests for this URL, even though they are otherwise blacklisted by the system for other URLs.

mobilesafe-encryption

Specifies, when enabled, that the system protects requests for this URL from mobile devices with Application layer encryption.

parameters

Adds, deletes, or replaces a set of sensitive parameters protected by the system. You can configure the following options for each parameter:

ajax-mapping

Specifies the mapping between the parameter name and its location in AJAX payload.

attach-to-vtoken-report

Specifies, when enabled, that the system adds the parameter value data to the alerts.

check-integrity

Specifies, when enabled, that the system verifies whether the user-input data was changed by malware when it left the browser for the server.

encrypt

Specifies, when enabled, that the system encrypts the parameter's value attribute.

identify-as-username

Specifies, when enabled, that the system considers this parameter a username. Note: There may be only one such parameter per URL, and its value is used only when login is successful (according to the URL's login-response).

method

Deprecated since v14.1.0. Please use parameter 'search-in' configuration instead. Specifies the method of the request from which the system gets the parameter data. Select either: POST (the default value) or GET.

mobilesafe-encrypt

Specifies that this parameter contains the encrypted fields from mobile devices. Note: There may be only one such parameter per URL (usually called auth), it cannot have other settings enabled and its method must be POST.

mobilesafe-entangle

Specifies that this parameter must be encrypted by mobile devices. The system replaces its value in the request payload and sends an alert if the mobilesafe-encrypt parameter does not contain this field.

obfuscate

Specifies, when enabled, that the system encrypts the parameter's name attribute.

priority

Specifies a unique ordinal number for this parameter in the set of wildcard parameters.

protect-by-selector

Specifies, when enabled, that the client considers this parameter's name to be a CSS selector. Note: To enable it, the parameter name must be defined as explicit and you must enable full-ajax-encryption.

search-in

Specifies the request part from which the system gets the parameter data. Select either: payload or query-string or any (the default value). If any is selected, then the query string will be searched first and only if the parameter is not found there, the payload will be also searched in.

substitute-value

Specifies, when enabled, that the system substitutes the parameter's value with asterisks [*] in the web application while the form is being filled. In order to enable it, you must first enable encrypt.

type Specifies a type of the parameter. Note: If you create a parameter, you can use either type, thereafter it becomes read only. The options are:

explicit

Specifies that the parameter has an exact path. This is the default value.

wildcard

Specifies that any parameter that matches this wildcard expression is considered protected.

phishing

Specifies when the system detects phishing attempts by attackers who set up a fake URL that imitates the real URL. You can configure the following options for Phishing detection:

capture-users

Specifies, when enabled, that the system logs the usernames and text fields (not passwords) of users attacked by a phishing attempt.

copy-detection

Specifies, when enabled, that the system detects copied web pages.

css-protection

Specifies, when enabled, that the system activates the CSS module, which is part of the system's phishing detection backup mechanism.

[enabled | disabled]

Specifies whether the system protects this URL against phishing, and sends an alert if the system detects this URL to be under a phishing attempt, or not.

field-types-to-send

Adds, deletes, or replaces a set of HTML input types whose values should be included in phishing alerts.

inject-css-element

Specifies where the system injects the CSS element. You can configure the following options for CSS element injection position:

[after | before]

Specifies whether the system injects the CSS element after an opening tag or before a closing tag.

tag Specifies the HTML tag for injection of the CSS element. This tag cannot be none.

inject-css-link

Specifies where the system injects the CSS link, when application-css is disabled. You can configure the following options for CSS link injection position:

[after | before]

Specifies whether the system injects the CSS link after an opening tag or before a closing tag.

tag Specifies the HTML tag for injection of the CSS link. This tag cannot be none.

inject-inline-javascript

Specifies where the system injects the phishing inline script and image. You can configure the following options for phishing inline script and image injection position:

[after | before]

Specifies whether the system injects the phishing inline script and image after an opening tag or before a closing tag.

tag Specifies the HTML tag for injection of the phishing inline script and image. This tag cannot be none.

priority

Specifies a unique ordinal number for this URL in the set of wildcard URLs.

type Specifies a type of the URL. Note: If you create a URL, you can use either type, thereafter it becomes read only. The options are:

explicit

Specifies that the URL has an exact path. This is the default value.

wildcard

Specifies that any URL that matches this wildcard expression is considered protected.

users

Adds, deletes, or replaces a set of users enforced by the system upon successful login. You can configure the following options for an enforced user:

modes

Adds or deletes a single mode in the set of existing user modes.

mode Specifies a unique mode for the user. This option is required for the operations add and delete. The options are:

block

Specifies that the system blocks the user account by displaying blocking-page.

forensic

Specifies that the system enforces the user to run Forensics tool on his host by displaying forensic.html.

inspection

Specifies that the system turns on verbose activity logging for this user, i.e. collects all HTML and JS sources from sessions and sends this data to the dashboard.

remediation

Specifies that the system enforces the user to run Forensics tool in remediation mode that deploys Anti-malware client on his host by displaying forensic.html.

duration

Specifies number of minutes during which the user is enforced in this mode since its first login, when enforce-policy is time-limited. After their expiration the user mode will be removed automatically.

enforce-policy

Specifies enforcement policy for this user mode. The options are:

enforce

Specifies that the user must download and run Forensics tool in order to continue online actions. Note: This policy may be specified only for the modes forensic and remediation.

time-limited

Specifies that the user is enforced in this mode for a limited time, namely until first-login-time + duration minutes. When this policy is specified for the modes forensic and remediation, the user may skip downloading and running Forensics tool every time.

unlimited

Specifies that the user is enforced in this mode for unlimited time. When this policy is specified for the modes forensic and remediation, the user may skip downloading and running Forensics tool every time.

first-login-time

Displays time when the user firstly logged in being in this mode. A new user mode is added with value none and it is updated automatically during traffic, when enforce-policy is time-limited.

whitelist-custom-alerts

Specifies a list of predefined alerts that are ignored.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, security, security anti-fraud, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2019-07-10 security anti-fraud profile(1)

security anti-fraud signatures-update

NAME

signatures-update - Runs or loads an Anti-fraud signatures update.

MODULE

security anti-fraud

SYNTAX

Configure, run or load the signatures-update component within the security anti-fraud module using the syntax in the following sections:

MODIFY

modify security anti-fraud signatures-update

options:
update-automatically [enabled | disabled]

edit security anti-fraud signatures-update

options:
all-properties
non-default-properties

LOAD

load signatures-update

options:
file [filename]

RUN

run signatures-update

DISPLAY

list signatures-update

options:
all-properties
current-version-create-datetime
download-available
install-datetime
install-user
message
non-default-properties
one-line
partition
progress-status
progress-status-datetime
last-update-check-datetime
readme
update-automatically

DESCRIPTION

You can use the signatures-update component to run, load, configure or display status of signatures update.

EXAMPLES

list security anti-fraud signatures-update

Displays the status of signatures update.

OPTIONS

current-version-create-datetime
Displays the creation time of currently installed signatures update version.

download-available
Displays whether new signatures version is available for download from the cloud.

file Specifies the file name from which the signatures update is going to be installed when using the load command. A full path should be specified.

install-datetime
Displays the time when signatures update was installed.

install-user
Displays the name of the user who installed the last signatures update.

message
Displays the message describing the failure status of signatures update.

partition
Displays the administrative partition within which this object resides.

progress-status
Displays the signatures update progress status.

progress-status-datetime
Displays the time when signatures update progress status was last changed.

last-update-check-datetime
Displays the time when last checked for signatures update.

readme
Displays the Readme content for the current signatures update.

update-automatically
Enables or disables automatic nightly update.

SEE ALSO

edit, list, modify, security, security anti-fraud, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

security blacklist-publisher all-blacklist-publisher

NAME

all-blacklist-publisher - Show command for blacklist publisher for use by firewall. An all-blacklist-publisher shows all the category records and their associated profiles.

MODULE

security blacklist-publisher

SYNTAX

Shows the category component within the security blacklist-publisher module using the syntax in the following sections.

show all-blacklist-publisher

EXAMPLES

show all-blacklist-publisher

Displays all the blacklist publisher records. Each record displays its category, profile association, route domain, traffic group, next hop, advertised count, and deadadvertised count.

FIELDS

category name
Specifies the category which is associated with the profile.

profile name
Specifies the profile which is associated with the category.

route domain
Specifies the route domain in the route advertisement.

traffic group
Specifies the traffic group which will be utilized in route injection fault tolerance.

next hop
Specifies the next-hop address which will be utilized in route advertisement in blackhole routes.

advertised count
Specifies the number of times this route has been advertised.

deadadvertised count
Specifies the number of times this route has been dropped.

SEE ALSO

show, profile category, security, security blacklist-publisher, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

security blacklist-publisher blacklist-publisher-stats

NAME

blacklist-publisher-stats - Show/reset-stats command for blacklist publisher for use by firewall. A blacklist-publisher-stats shows all the category records and their associated profiles.

MODULE

security blacklist-publisher

SYNTAX

Shows the category component within the security blacklist-publisher module using the syntax in the following sections.

```
show blacklist-publisher-stats
show blacklist-publisher-stats category-name [name]
show blacklist-publisher-stats profile-name [name]
show blacklist-publisher-stats category-name [name] profile-name [name]
```

Reset-Stats the category component within the security blacklist-publisher module using the syntax in the following sections.

```
reset-stats blacklist-publisher-stats
reset-stats blacklist-publisher-stats category-name [name]
reset-stats blacklist-publisher-stats profile-name [name]
reset-stats blacklist-publisher-stats category-name [name] profile-name [name]
```

SHOW EXAMPLES

```
show blacklist-publisher-stats
```

Displays all the blacklist publisher records. Each record displays its category, profile association, advertised count, and deadadvertised count.

FIELDS

`category_name`
Specifies the category which is associated with the profile.

`profile_name`
Specifies the profile which is associated with the category.

`advertise_it`
Specifies the number of times this route has been advertised.

`deadvertise_it`
Specifies the number of times this route has been dropped.

RESET_STATS EXAMPLES

```
reset-stats blacklist-publisher-stats
```

Reset stats for all the blacklist publisher records.

```
reset-stats blacklist-publisher-stats category_name <\Common\botnets
```

Reset stats of all blacklist publisher records that are associated with "\Common\botnets" category.

```
reset-stats blacklist-publisher-stats profile_name <\Common\botnets-profile
```

Reset stats of all blacklist publisher records that are associated with "\Common\botnets-profile" profile.

```
reset-stats blacklist-publisher-stats category_name <\Common\botnets profile_name <\Common\botnets-profile>
```

Reset stat of all the blacklist publisher records that are associated with "\Common\botnets" category and "\Common\botnets-profile" profile.

SEE ALSO

show, reset-stats, profile category, security, security blacklist-publisher, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2017-1security blacklist-publisher blacklist-publisher-stats(1)

security blacklist-publisher by-addr

NAME

by-addr - Show command for blacklist publisher for use by firewall. A by-addr shows all the category records and their associated profiles for the specified shun address.

MODULE

security blacklist-publisher

SYNTAX

Shows the shun-ip component within the security blacklist-publisher module using the syntax in the following sections.

```
show by-addr [enter IP address, IP address/prefix, or IP address range (two addresses separated by a hyphen).]
```

EXAMPLES

```
show by-addr 8.8.8.8
```

show by-addr 8.8.8.8/24

show by-addr 8.8.8.8-8.8.8.10

Displays all the blacklist publisher records for the specified shun address.

FIELDS

IP Address

Specifies the shun ip to be advertised.

category name

Specifies the category which is associated with the profile.

profile name

Specifies the profile which is associated with the category.

route domain

Specifies the route domain in the route advertisement.

traffic group

Specifies the traffic group which will be utilized in route injection fault tolerance.

next hop

Specifies the next-hop address which will be utilized in route advertisement in blackhole routes.

time-to-live

Specifies the amount of time left until this route will be dropped.

SEE ALSO

show, profile category, security, security blacklist-publisher, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2017-10-24 security blacklist-publisher by-addr(1)

security blacklist-publisher by-category

NAME

by-category - Show command for blacklist publisher for use by firewall. A by-category shows all the category records and their associated profiles for the specified category name.

MODULE

security blacklist-publisher

SYNTAX

Shows the category component within the security blacklist-publisher module using the syntax in the following sections.

show by-category category-name [[name] | all]

EXAMPLES

show by-category category-name /Common/botnets show by-category category-name all

Displays all the blacklist publisher records for the specified category.

FIELDS

category name

Specifies the category which is associated with the profile.

profile name

Specifies the profile which is associated with the category.

route domain

Specifies the route domain in the route advertisement.

traffic group

Specifies the traffic group which will be utilized in route injection fault tolerance.

next hop

Specifies the next-hop address which will be utilized in route advertisement in blackhole routes.

shun-ip

Specifies the shun ip to be advertised.

time-to-live

Specifies the amount of time left until this route will be dropped.

advertised

Specifies whether this record is currently being advertised.

SEE ALSO

show, profile category, security, security blacklist-publisher, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2015-08-28 security blacklist-publisher by-category(1)

security blacklist-publisher category

NAME

category - Configures a blacklist publisher category for use by firewall. A category enables a user to advertise shuns based on the route-advertisement-next-hop value(s) in the category's associated profile(s).

MODULE

security blacklist-publisher

SYNTAX

Configure the category component within the security blacklist-publisher module using the syntax in the following sections.

CREATE/MODIFY

create category [name]

modify category [name]

options:

profile-names [add | default | delete | none | replace-all-with] {
[name]
}

edit category [[name] | all]

options:

all-properties
non-default-properties

list category [[name] | all | [property]]

show running-config profile [[name] | all | [property]]

options:

all-properties
non-default-properties
one-line

DELETE

delete category [[name] | all | recursive]

Note: If the category is referencing any profile the delete will fail.

DESCRIPTION

You can use the category component to create, modify, display, or delete a blacklist-publisher category.

EXAMPLES

```
create category botnets profile-names add { botnets-profile }
```

Creates the above category and adds the profile (botnets-profile).

```
modify category botnets profile-names delete { botnets-profile }
```

Modifies the above category by deleting the profile (botnets-profile).

```
delete category botnets
```

Deletes the above category.

```
list category botnets
```

Displays the properties of the above category.

OPTIONS

profile-names

Specifies the profiles referenced in the category.

SEE ALSO

create, delete, edit, modify, list, profile, security, security blacklist-publisher, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2015-08-28 security blacklist-publisher category(1)

security blacklist-publisher profile

NAME

profile - Configures a publisher profile for use by firewall. A profile specifies the network path and method (how and where) the IP-Intelligence Blacklist data will be advertised.

MODULE

security blacklist-publisher

SYNTAX

Configure the profile component within the security blacklist-publisher module using the syntax in the following sections.

CREATE/MODIFY

create profile [name]

modify profile [name | all]

options:

bgp-flowspec-advertisement-action [drop | rate-limit | qos]

bgp-flowspec-dscp-value [integer]

bgp-flowspec-rate-limit [integer]

description [string]

route-domain [name]

traffic-group [name]

route-advertisement-nextthop [IP address]

route-advertisement-nextthop-v6 [IP address]

edit profile [[name] | all]

options:

all-properties

non-default-properties

list profile [[name] | all | [property]]

show running-config profile [[name] | all | [property]]

options:

all-properties

non-default-properties

one-line

DELETE

delete profile [[name] | all | recursive]

Note: If the profile is referenced in a category the delete will fail.

DESCRIPTION

You can use the profile component to create, modify, display, or delete a blacklist-publisher profile.

EXAMPLES

```
create profile botnets-profile { route-domain 0 route-advertisement-nextthop 10.10.10.1 }
```

Creates the above profile with the specified route-domain, and route-advertisement-nextthop.

```
modify profile botnets-profile description "this profile is used for botnets category"
```

Modifies the above profile with a description.

```
delete profile botnets-profile
```

Deletes the above profile.

```
list profile botnets-profile
```

Displays the properties of the above profile.

OPTIONS

bgp-flowspec-advertisement-action

Specifies the BGP FlowSpec Advertisement Action to be used for Blackholing. The default is drop.

`bgp-flowspec-dscp-value`
Specifies the BGP FlowSpec DSCP value for advertisement qos action.

`bgp-flowspec-rate-limit`
Specifies the BGP FlowSpec rate limit (bytes/sec) for advertisement rate limiting action.

`description`
User defined description.

`route-advertisement-nexthop`
Specifies the next-hop v4 address which will be utilized in route advertisement in blackhole routes.

`route-advertisement-nexthop-v6`
Specifies the next-hop v6 address which will be utilized in route advertisement in blackhole routes.

`route-domain`
Specifies the route-domain in the route advertisement.

`traffic-group`
Deprecated since v13.1.0. Specifies the traffic group which will be utilized in route injection fault tolerance.

SEE ALSO

`create`, `delete`, `edit`, `list`, `modify`, `category`, `security`, `security blacklist-publisher`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2015. All rights reserved.

BIG-IP 2018-06-08 security blacklist-publisher profile(1)

security bot-defense anomaly-category

NAME

`anomaly-category` - Shows the Bot Defense anomalies categories.

MODULE

`security bot-defense`

SYNTAX

Shows the `anomaly-category` component within the `security bot-defense` module using the syntax shown in the following sections.

DISPLAY

`list anomaly-category`

DESCRIPTION

You can use the `anomaly-category` component to display a Bot Defense Anomaly Category

EXAMPLES

`list anomaly-category`

Displays the properties of all Bot Defense Anomalies Categories.

OPTIONS

`class`

The class of the category.

`description`

A description of the category.

SEE ALSO

`list`, `security`, `security bot-defense`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2018. All rights reserved.

BIG-IP 2018-10-20 security bot-defense anomaly-category(1)

security bot-defense anomaly

NAME

anomaly - Shows the Bot Defense anomalies.

MODULE

security bot-defense

SYNTAX

Shows the anomaly component within the security bot-defense module using the syntax shown in the following sections.

DISPLAY

list anomaly

DESCRIPTION

You can use the anomaly component to display a Bot Defense Anomaly

EXAMPLES

list anomaly

Displays the properties of all Bot Defense Anomalies.

OPTIONS

type Specifies the anomaly type. The possible values are general, js-free, js-verification, cshui, micro-service.

category

The category of the anomaly.

description

A description of the anomaly.

SEE ALSO

list, security, security bot-defense, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2018. All rights reserved.

BIG-IP 2018-10-20 security bot-defense anomaly(1)

security bot-defense class

NAME

class - Shows the bot defense classes

MODULE

security bot-defense

SYNTAX

Shows the class component within the security bot-defense module using the syntax shown in the following sections.

DISPLAY

list class

DESCRIPTION

You can use the class component to display a Bot Defense Class

EXAMPLES

list class

Displays the properties of all Bot Defense Classes.

OPTIONS

description

A description of the anomaly.

micro-services

The default mitigation or verification action for this class in each micro service.

templates

The default mitigation or verification action for this class in each template.

type The type of the class. The options are: none, browser, mobile, trusted-bot, untrusted-bot, malicious-bot, suspicious-browser, unknown.

SEE ALSO

list, security, security bot-defense, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2018. All rights reserved.

BIG-IP 2018-10-20 security bot-defense class(1)

security bot-defense micro-service

NAME

micro-service - Shows the Bot Defense Microservices Types

MODULE

security bot-defense

SYNTAX

Shows the micro-service component within the security bot-defense module using the syntax shown in the following sections.

DISPLAY

list bot-defense-class

DESCRIPTION

You can use the micro-service component to display a Bot Defense Microservices Types.

EXAMPLES

list micro-service

Displays the properties of all Bot Defense Microservices Types.

OPTIONS

detection-threshold

The default number of times the microservice was accessed by a client after which the microservice anomaly is declared, for this type.

detection-time

The default period in seconds for this type during which the microservice anomaly threshold is tested.

intent

The corresponding OWASP threat.

mitigation-time

The time in seconds for the microservice anomaly mitigation duration, for this type.

SEE ALSO

list, security, security bot-defense, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2018. All rights reserved.

BIG-IP 2018-10-20 security bot-defense micro-service(1)

security bot-defense profile

NAME

profile - Configures a Bot Defense profile.

MODULE

security bot defense

SYNTAX

Configure the profile component within the security bot defense module using the syntax shown in the following sections.

CREATE/MODIFY

create profile [name]

modify profile [name]

options:

api-access-strict-mitigation [enabled | disabled]

anomaly-category-overrides [none | add | delete | modify | replace-all-with] {
[anomaly-category-name] ... }

options:

action [alarm | block | captcha | none | tcp-reset | redirect-to-pool | honeypot-page]

app-service [[string] | none]

}

}

anomaly-overrides [none | add | delete | modify | replace-all-with] {

[anomaly-name] }

options:

action [alarm | block | captcha | none | tcp-reset | redirect-to-pool | honeypot-page]

}

}

app-service [[string] | none]

blocking-page {

body [string]

headers [string]

status-code [integer]

type [custom | default]

}

captcha-response {

failure {

body [string]

type [custom | default]

}

first {

body [string]

type [custom | default]

}

class-overrides [none | add | delete | modify | replace-all-with] {

[class-name] }

options:

mitigation {

action [alarm | block | captcha | none | rate-limit | tcp-reset | redirect-to-pool | honeypot-page]

rate-limit-tps [integer]

}

verification {

action [browser-challenge-free-verification | browser-verify-after-access-detection | mobile-verify-integrity browser-verify-after-access-bloc

}

}

cross-domain-requests [allow-all | validate-bulk | validate-upon-request]

defaults-from [profile-name]

description [[string] | none]

deviceid-mode [generate-after-access | generate-before-access | none]

dos-attack-strict-mitigation [enabled | disabled]

enforcement-mode [blocking | transparent]

enforcement-readiness-period [integer]

external-domains [none | add | delete | modify | replace-all-with] { [string] ... }

grace-period [integer]

micro-services [none | add | delete | modify | replace-all-with] {

action [alarm | block | captcha | none | tcp-reset | redirect-to-pool | honeypot-page]

class-overrides [none | add | delete | modify | replace-all-with] {

[class-name] }

options:

mitigation {

action [alarm | block | captcha | none | tcp-reset | redirect-to-pool | honeypot-page]

}

verification {

action [browser-captcha | browser-challenge-free-verification | browser-verify-after-access-detection | mobile-verify-integrity browser-verif

}

}

}

description [[string] | none]

detection-threshold [integer]

detection-time [integer]

enforcement-mode [blocking | profile-default | transparent]

hostname [string]

match-order [integer]

mitigation-time [integer]

type [micro-service-type]

urls [none | add | delete | modify | replace-all-with] {

match-order [integer]

url [string]

```

}
}
mobile-detection {
  allow-android-rooted-device [enabled | disabled]
  allow-any-android-package [enabled | disabled]
  allow-any-ios-package [enabled | disabled]
  allow-emulators [enabled | disabled]
  allow-jailbroken-devices [enabled | disabled]
  android-publishers [none | add | delete | modify | replace-all-with] { [string] ... }
  block-debugger-enabled-device [enabled | disabled]
  client-side-challenge-mode [cshui | pass]
  ios-allowed-packages [none | add | delete | modify | replace-all-with] { [string] ... }
  signatures [none | add | delete | modify | replace-all-with] { [signature-name] ... }
}
perform-challenge-in-transparent [enabled | disabled]
signature-category-overrides [none | add | delete | modify | replace-all-with] {
  [signature-category] ... {
    options:
  }
}
action [alarm | block | captcha | none | tcp-reset | redirect-to-pool | honeypot-page]
}
}
signature-overrides [none | add | delete | modify | replace-all-with] {
  [signature-name] ... {
    options:
  }
}
action [alarm | block | captcha | none | rate-limit | tcp-reset | redirect-to-pool | honeypot-page]
rate-limit-tps [integer]
}
}
signature-staging-upon-update [enabled | disabled]
single-page-application [enabled | disabled]
site-domains [none | add | delete | modify | replace-all-with] { [string] ... }
staged-signatures [none | add | delete | modify | replace-all-with] { [signature-name] ... }
template [balanced | relaxed | strict]
whitelist [none | add | delete | modify | replace-all-with] {
  disable-mitigation [no | yes]
  disable-verification [no | yes]
  geolocation [country-name]
  match-order [integer]
  source-address [ip-address]
  url [string]
}
}

```

```

edit profile [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

```

```

DISPLAY
list profile
list profile [ [ [name] | [glob] | [regex] ] ... ]
show running-config profile
show running-config profile [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  recursive
show profile [ [ [name] | [glob] | [regex] ] ... ]

```

```

DELETE
delete profile [name]

```

DESCRIPTION

You can use the profile component to create, modify, display, or delete a Bot Defense profile for use with Bot Defense Protection functionality.

EXAMPLES

```
create profile my_bot_profile
```

Creates a custom Bot Defense profile named my_bot_profile with initial settings.

```
list profile
```

Displays the properties of all Bot Defense profiles.

OPTIONS

api-access-strict-mitigation
Determines, when enabled, whether to apply the strict mitigation settings on API access requests.

anomaly-category-overrides
List or define anomaly categories that have mitigation settings overriding the defaults determined by the template.

action
The overriding mitigation action for the anomaly category. The options are:

alarm

Mark the request log entry with the alarm flag. Request passes through.

block

Block the request and send the blocking page in response.

captcha

Send a CAPTCHA challenge in response. If the user solves the challenge successfully the requests from this user are passed to the server for a limited period of time.

none The category will not be detected. This may affect the classification of the client and hence the action based on that classification.

tcp-reset

Discard the request by resetting the TCP connection it was sent on.

redirect-to-pool

Send the request to the configured pool.

honeypot-page

Block the request and send the Honeypot response page in response.

anomaly-overrides

List or define anomalies that have mitigation settings overriding the defaults determined by the template.

action

The overriding mitigation action for the anomaly. The options are:

alarm

Mark the request log entry with the alarm flag. Request passes through.

block

Block the request and send the blocking page in response.

captcha

Send a CAPTCHA challenge in response. If the user solves the challenge successfully the requests from this user are passed to the server for a limited period of time.

none The anomaly will not be detected. This may affect the classification of the client and hence the action based on that classification.

tcp-reset

Discard the request by resetting the TCP connection it was sent on.

redirect-to-pool

Send the request to the configured pool.

honeypot-page

Block the request and send the Honeypot response page in response.

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

blocking-page

body Configures the body of the blocking page response.

headers

Configures the HTTP headers of the blocking page response.

status-code

Configures the HTTP status code of the blocking page response.

type Configures a type of the blocking page response.

captcha-response

Specifies properties of the CAPTCHA response in Application Security. You can configure the following options for CAPTCHA Response Settings:

failure

Specifies properties of a failed CAPTCHA response. You can configure the following options for a failed CAPTCHA response:

body Configures a failed CAPTCHA response body.

type Configures a type of a failed CAPTCHA response body. You can configure the following options for a failed CAPTCHA response type:

custom

Configures a custom failed CAPTCHA response type.

default

Configures a default failed CAPTCHA response type.

first

Specifies properties of the first CAPTCHA response. You can configure the following options for the first CAPTCHA response:

body Configures the first CAPTCHA response body.

type Configures a type of the first CAPTCHA response body. You can configure the following options for the first CAPTCHA response type:

custom

Configures a custom first CAPTCHA response type.

default

Configures a default first CAPTCHA response type.

class-overrides

List or define classes that have settings overriding the defaults determined by the template.

mitigation

These set which type of mitigation to take for each type of bot.

action

The overriding mitigation action for the class. The options are:

alarm

Mark the request log entry with the alarm flag. Request passes through.

block

Block the request and send the blocking page in response.

captcha

Send a CAPTCHA challenge in response. If the user solves the challenge successfully the requests from this user are passed to the server for a limited period of time.

none The class will not be detected. This may affect the classification of the client and hence the action based on that classification.

rate-limit

Traffic from this source will be limited to a given transaction rate. Relevant only to classes and signature categories.

tcp-reset

Discard the request by resetting the TCP connection it was sent on.

redirect-to-pool

Send the request to the configured pool.

honeypot-page

Block the request and send the Honeypot response page in response.

rate-limit-tps

The TPS limit for the signature Relevant only if the action is rate limit.

verification

action

The overridden verification action for the class. Relevant only to Browser and Mobile Application classes. The options are:

browser-challenge-free-verification

Verify the browser authenticity by examining the request header without sending any JavaScript challenge to the client.

browser-verify-after-access-detection

Verify the browser authenticity by examining the request header and by a JavaScript challenge injected to the server response after letting the request to the server. The request is not mitigated even if the JavaScript challenge failed, rather, it is only reported, classifying the client as a Bot.

browser-verify-after-access-blocking

Verify the browser authenticity by examining the request header and by a JavaScript challenge injected to the server response after letting the request to the server. Upon challenge failure, the request is mitigated according to the respective mitigation settings.

browser-verify-before-access

Verify the browser authenticity by examining the request header and by responding with a JavaScript challenge before letting the request to the server. Request is sent to the server only upon successful completion of the challenge.

mobile-verify-integrity

Enable verification of mobile applications with Anti-Bot mobile SDK.

none No verification action taken. This may affect the classification of the client and hence the action based on that classification.

cross-domain-requests

Specifies a cross-domain requests handling mode. The options are:

`allow-all`
Allows all cross-domain requests. This is the default value.

`validate-bulk`
System validates domains in bulk: the cookies for the related domains are created together with the cookie for the current domain, by generating challenges in iframes - one per each domain.

`validate-upon-request`
System validates domains upon request: the cookie for the related domain is generated when a request arrives to an unqualified URL without a cookie.

`defaults-from`
Specifies the profile that you want to use as the parent profile.

`description`
User textual description of the profile.

`deviceid-mode`
Specifies how the Device ID will be collected.

The options are:

`generate-after-access`
Device ID is collected by injecting a JavaScript in the response. This is less intrusive and has less of a latency impact. Use when sensitive to response time or time-to-first-byte. Note that some requests at the beginning of the session will have no Device ID until the handshake is complete and the Device ID is generated.

`generate-before-access`
Device ID is collected by generating a JavaScript challenge before forwarding the HTTP request to the application. This guarantees that every request that reaches the application has a Device ID. This has more of a latency impact than Generate After Access. Bots that present as browsers and are unable to execute JavaScript will be blocked.

`none` Device ID will not be collected.

`dos-attack-strict-mitigation`
Determines, when enabled, whether to apply strict mitigation settings during DoS attacks.

`enforcement-mode`
Specifies whether to execute the mitigation actions or just report the detection results.

`blocking`
The system will replace responses with client side JavaScript and if the client cannot run JavaScript, it will not be able to receive the server responses.

`transparent`
The system logs all mitigation and verification actions but no action is taken on the traffic.

`enforcement-readiness-period`
How many days, since the bot defense profile was last changed, that the profile remains in staging mode before the system suggests you enforce. The system does not enforce profile entities and attack signatures in staging. Staging allows you to test the bot defense profile entities and the attack signatures for false positives without enforcing them. The default is 604800 seconds (7 days).

`external-domains`
List of external domain names that may refer to resources on the application. Used for handling cross-domain requests.

`grace-period`
The period in seconds during which requests for resources are not required to have browser verification cookies.

`micro-services`
List of microservices for which there exists specialized detection and mitigation behavior.

`action`
The mitigation action for the anomaly of the microservice. The options are:

`alarm`
Mark the request log entry with the alarm flag. Request passes through.

`block`
Block the request and send the blocking page in response.

`captcha`
Send a CAPTCHA challenge in response. If the user solves the challenge successfully the requests from this user are passed to the server for a limited period of time.

`none` Microservice anomaly will not be detected. This may affect the classification of the client and its mitigation action.

`tcp-reset`
Discard the request by resetting the TCP connection it was sent on.

`redirect-to-pool`
Send the request to the configured pool.

honeypot-page

Block the request and send the Honeypot response page in response.

class-overrides

The list of classes for which the settings override the ones determined by the microservice type.

mitigation

These set which type of mitigation to take for each type of bot.

action

The name of the overriding mitigation action for the class in the scope of the microservice. The options are:

alarm

Mark the request log entry with the alarm flag. Request passes through.

block

Block the request and send the blocking page in response.

captcha

Send a CAPTCHA challenge in response. If the user solves the challenge successfully the requests from this user are passed to the server for a limited period of time.

none The class will not be detected. This may affect the classification of the client and hence the action based on that classification.

rate-limit

Traffic from this source will be limited to a given transaction rate. Relevant only to classes and signature categories.

tcp-reset

Discard the request by resetting the TCP connection it was sent on.

redirect-to-pool

Send the request to the configured pool.

honeypot-page

Block the request and send the Honeypot response page in response.

verification

action

The name of the overriding verification action for the class in the scope of the microservice. Relevant only to Browser class. The options are:

browser-captcha

Send a CAPTCHA challenge response on each request with Browser User-Agent. The request is passed to the server only if the CAPTCHA was successfully solved.

browser-challenge-free-verification

Verify the browser authenticity by examining the request header without sending any JavaScript challenge to the client.

browser-verify-after-access-detection

Verify the browser authenticity by examining the request header and by a JavaScript challenge injected to the server response after letting the request to the server. The request is not mitigated even if the JavaScript challenge failed, rather, it is only reported, classifying the client as a Bot.

browser-verify-after-access-blocking

Verify the browser authenticity by examining the request header and by a JavaScript challenge injected to the server response after letting the request to the server. Upon challenge failure, the request is mitigated according to the respective mitigation settings.

browser-verify-before-access

Verify the browser authenticity by examining the request header and by responding with a JavaScript challenge before letting the request to the server. Request is sent to the server only upon successful completion of the challenge.

none No verification action taken. This may affect the classification of the client and hence the action based on that classification.

description

User textual description of the microservice.

detection-threshold

The number of times the microservice was accessed by a client after which the microservice anomaly is declared.

detection-time

The period in seconds during which the microservice anomaly threshold is tested.

enforcement-mode

Determines the enforcement mode for the microservice, or use the default one from the profile settings.

blocking

The system will replace responses with client side JavaScript and if the client cannot run JavaScript, it will not be able to receive the server responses.

profile-default

The enforcement mode (transparent or blocking) for the microservice is determined by the profile setting.

transparent

The system logs all mitigation and verification actions but no action is taken on the traffic.

hostname

The host domain name used to access the microservice. Wildcards are supported.

match-order

Ordinal number (1,2,3...) specifying the order that the microservice entries are checked for matches.

mitigation-time

The time in seconds for the microservice anomaly mitigation duration.

type The type of the microservice.

urls

match-order

Ordinal number (1,2,3...) specifying the order that the URL entries in the microservice are checked for matches.

url The URL entry in the microservice. Wildcards are supported.

mobile-detection

This feature detects mobile applications built with the Anti-Bot Mobile SDK and defines how requests from these mobile application clients are handled. This feature requires an Anti-Bot Mobile SK license to be operational.

allow-android-rooted-device

Specifies, when enabled, whether to allow mobile applications from Android rooted devices.

allow-any-android-package

Specifies, when enabled, whether to allow any mobile applications from Android devices.

allow-any-ios-package

Specifies, when enabled, whether to allow any mobile applications from iOS devices.

allow-emulators

Specifies, when enabled, whether to allow mobile applications running on emulators.

allow-jailbroken-devices

Specifies, when enabled, whether to allow mobile applications from iOS jailbroken devices.

android-publishers

List of allowed publisher certificates for Android mobile applications.

block-debugger-enabled-device

Specifies, when enabled, whether to allow mobile applications from devices with debugger enabled.

client-side-challenge-mode

Determines the substitute action for CAPTCHA or client-side integrity check for mobile applications. In those cases, the options for the mobile applications are:

cshui

The SDK checks for human interactions with the screen in the last few seconds. If none are detected, the traffic is blocked.

pass The traffic is passed without incident.

ios-allowed-packages

List of fully qualified iOS mobile application package names that will be allowed.

signatures

List of signature names for detecting mobile applications without Anti-Bot Mobile SDK.

perform-challenge-in-transparent

Determines, when enabled, whether browser challenges are still executed in transparent enforcement mode.

signature-category-overrides

List of signature categories that have mitigation settings overriding the defaults determined by the template.

action

The overriding mitigation action for the Bot signature category. The options are:

alarm

Mark the request log entry with the alarm flag. Request passes through.

block

Block the request and send the blocking page in response.

captcha

Send a CAPTCHA challenge in response. If the user solves the challenge successfully the requests from this user are passed to the server for a limited period of time.

none The category will not be detected. This may affect the classification of the client and hence the action based on that classification.

tcp-reset

Discard the request by resetting the TCP connection it was sent on.

redirect-to-pool

Send the request to the configured pool.

honeypot-page

Block the request and send the Honeypot response page in response.

signature-overrides

List of signatures that have mitigation settings overriding the defaults determined by the template.

action

The overriding mitigation action for the Bot signature. The options are:

alarm

Mark the request log entry with the alarm flag. Request passes through.

block

Block the request and send the blocking page in response.

captcha

Send a CAPTCHA challenge in response. If the user solves the challenge successfully the requests from this user are passed to the server for a limited period of time.

none The signature will not be detected. This may affect the classification of the client and hence the action based on that classification.

rate-limit

Traffic from this source will be limited to a given transaction rate. Relevant only to classes and signature categories.

tcp-reset

Discard the request by resetting the TCP connection it was sent on.

redirect-to-pool

Send the request to the configured pool.

honeypot-page

Block the request and send the Honeypot response page in response.

rate-limit-tps

The TPS limit for the signature Relevant only if the action is rate limit.

signature-staging-upon-update

In staging, the system does not block the request, but logs the request. When not in staging, the system enforces the mitigation action configured for the signature in Mitigation Settings.

single-page-application

Enable if your website is a Single Page Application, meaning a web application that loads new content without triggering a full page-reload. The system will inject JavaScript code to every HTML response. This will allow handling browser challenges and CAPTCHA without requiring page reloading. The default is Disabled.

site-domains

Configures a list of domains that are part of the website.

staged-signatures

List of signatures that are put in staging.

template

Specifies the template for the profile. The template determines default values for the mitigation settings and some other settings. The options are:

balanced

Allow limited access to non-malicious bots and verify browsers without affecting the user experience.

relaxed

Allow full access to non-malicious bots and perform non-intrusive browser verification.

strict

Allow only fully verified browsers, mobile applications and trusted bots and block the rest. User experience may be affected.

whitelist

These are the whitelisted sources. A source can be an IP address or a geolocation.

disable-mitigation

Specifies whether to exempt the whitelist entry from mitigation actions.

`disable-verification`
Specifies whether to exempt the whitelist entry from browser verification actions.

`geolocation`
The name of the source geolocation that is whitelisted.

`match-order`
Ordinal number (1,2,3...) specifying the order that the whitelist entries are checked for matches.

`source-address`
The name of the source IP address that is whitelisted.

`url` The name of the URL (wildcard supported) that is whitelisted.

`glob` Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

`partition`
Displays the administrative partition within which the component resides.

`regex`
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

SEE ALSO

`create`, `delete`, `edit`, `glob`, `list`, `ltm virtual`, `modify`, `regex`, `security`, `security bot-defense`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015, 2018. All rights reserved.

BIG-IP 2019-02-10 security bot-defense profile(1)

security bot-defense signature-category

NAME
signature-category - Configures the Bot Defense Signature Categories.

MODULE
security dos

SYNTAX
Configure the signature-category component within the security bot-defense module using the syntax shown in the following sections.

CREATE/MODIFY
create signature-category [string]
modify signature-category [name]
options:
class [class-name]

DISPLAY
list signature-category

DELETE
delete signature-category [name]

DESCRIPTION
You can use the signature-category component to create, modify, display, or delete a Bot Defense Signature Category.

EXAMPLES
create signature-category my_signature_category

Creates a custom Bot Defense Signature Category named my_signature_category with initial settings.

list signature-category

Displays the properties of all Bot Defense Signature Categories.

OPTIONS

class

Specifies the class to which the category belongs.

SEE ALSO

edit, list, modify, security, security dos, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2018. All rights reserved.

BIG-IP 2018-10-20 security bot-defense signature-category(1)

security bot-defense signature

NAME

signature - Configures the Bot Defense Signatures.

MODULE

security dos

SYNTAX

Configure the signature component within the security bot-defense module using the syntax shown in the following sections.

CREATE/MODIFY

create signature [string]

modify signature [name]

options:

category [name]

domains [none | add | delete | replace-all-with] { [string] ... }

risk [high | low | medium]

rule [string]

signature-references [string]

url {

match-type [contains | regexp]

search-string [string]

}

user-agent {

match-type [contains | regexp]

search-string [string]

}

DISPLAY

list signature

DELETE

delete signature [name]

DESCRIPTION

You can use the signature component to create, modify, display, or delete a Bot Defense Signature.

EXAMPLES

```
create signature my_signature
```

Creates a custom Bot Defense Signature named my_signature with initial settings.

```
list signature
```

Displays the properties of all Bot Defense Signatures.

OPTIONS

category

Specifies the Bot category to which the Bot signature belongs.

domains

Specifies the domain names from which the client has to come for the respective signature. Required for signatures of the Trusted Bot class, optional for others

risk Specifies the risk from Bot detected by the signature. The possible values are high, low and medium.

rule Specifies the signature matching rule in Snort-like format. Advanced alternative to the User-Agent/URL search strings.

signature-references

Specifies references to resources on the signature.

url Specifies the bot signature's URL matching rule. The following options are available:

match-type

Specifies how the URL search string is matched: simple substring or regexp match. The possible values are contains or regexp.

search-string

Specifies the string to be matched to the URL string in order for the signature to be detected.

user-agent

Specifies the bot signature's User-Agent matching rule. The following options are available:

match-type

Specifies how the User-Agent search string is matched: simple substring or regexp match. The possible values are contains or regexp.

search-string

Specifies the string to be matched to the User-Agent string in order for the signature to be detected.

SEE ALSO

edit, list, modify, security, security dos, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2018. All rights reserved.

BIG-IP 2018-10-20 security bot-defense signature(1)

security bot-defense template

NAME

template - Shows the template for the profile. The template determines default values for the mitigation settings and some other settings

MODULE

security bot-defense

SYNTAX

Shows the template component within the security bot-defense module using the syntax shown in the following sections.

DISPLAY

list template

DESCRIPTION

You can use the template component to display a Bot Defense Template

EXAMPLES

list template

Displays the properties of all Bot Defense Templates

OPTIONS

api-access-strict-mitigation

Determines, when enabled, whether to apply the strict mitigation settings on API access requests in this template.

deviceid-mode

Specifies how the Device ID will be collected in this template.

The options are:

generate-after-access

Device ID is collected by injecting a JavaScript in the response. This is less intrusive and has less of a latency impact. Use when sensitive to response time or time-to-first-byte. Note that some requests at the beginning of the session will have no Device ID until the handshake is complete and the Device ID is generated.

generate-before-access

Device ID is collected by generating a JavaScript challenge before forwarding the HTTP request to the application. This guarantees that every request that reaches the application has a Device ID. This has more of a latency impact than Generate After Access. Bots that present as browsers and are unable to execute JavaScript will be blocked.

none
dos-attack-strict-mitigation
Determines, when enabled, whether to apply strict mitigation settings during DoS attacks in this template.

SEE ALSO

list, security, security bot-defense, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2018. All rights reserved.

BIG-IP 2018-10-20 security bot-defense template(1)

security cloud-services cmd

NAME

cmd - Commands to the Security Cloud Services

MODULE

security cloud-services

SYNTAX

Send commands to the Security Cloud Services using the syntax shown in the following sections.

RUN

run cmd activate start
run cmd activate finalize
run cmd reactivate start
run cmd reactivate finalize
run cmd reload connector
run cmd portal open

DESCRIPTION

You can use cmd to send various commands to the Security Cloud Services component.

OPTIONS

activate

Activate the Cloud Services for the first time. The command will fail if activation is complete. First call start and a login link will be returned. Open the link in a browser, and follow the instructions within. Then, call finalize to complete the process.

start

Start the activation process. A login link will be returned. Open the link in a browser, and follow the instructions within.

finalize

Finalize the activation process.

reactivate

Re-activate the Cloud Services. The command will fail if activation was not previously done. First call start and a login link will be returned. Open the link in a browser, and follow the instructions within. Then, call finalize to complete the process.

start

Start the re-activation process. A login link will be returned. Open the link in a browser, and follow the instructions within.

finalize

Finalize the re-activation process.

reload connector

Reload the connector details from the Cloud Services.

portal open

Open the Cloud Services Portal. A temporary login link will be returned which must be opened in a browser.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

security cloud-services connector

NAME

connector - Configure the Cloud Services Connector

MODULE

security cloud-services

SYNTAX

Manage the connector component within the security cloud-services module using the syntax shown in the following sections.

CREATE/MODIFY

create connector [string]

modify connector [name]

options:

activation-time [time]

activation-note [string]

expiration-time [time]

deployment-id [string]

description [string]

params [string]

control-url [string]

control-token [string]

control-key [string]

clientside-url [string]

clientside-token [string]

clientside-key [string]

services {

centralized-device-id {

[enabled | disabled]

}

}

edit connector [name]

options:

all-properties

non-default-properties

DISPLAY

list connector

list connector [name]

options:

all-properties

non-default-properties

one-line

DELETE

delete connector [name]

DESCRIPTION

You can use the connector component to create, modify, display, or delete a Cloud Services connector.

OPTIONS

description

User defined description.

activation-time

Specifies the activation time of the connector.

activation-note

Specifies the activation note of the connector.

expiration-time

Specifies the expiration time of the connector.

deployment-id

Specifies the deployment id of the connector.

params

Specifies the custom parameters of the connector.

control-url

Specifies the control url of the connector.

control-token
Specifies the control token of the connector.

control-key
Specifies the control key of the connector.

clientside-url
Specifies the clientside url of the connector.

clientside-token
Specifies the clientside token of the connector.

clientside-key
Specifies the clientside key of the connector.

services
You can configure the following services:

centralized-device-id
You can configure the following options for Centralized Device ID:

enabled
Enables the Centralized Device ID service.

disabled
Disables the Centralized Device ID service.

SEE ALSO
create, edit, list, modify, delete, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2018. All rights reserved.

BIG-IP 2018-05-29 security cloud-services connector(1)

security datasync background-tasks

NAME
background-tasks - Configure the Schedule of the Datasync Background Tasks

MODULE
security datasync

SYNTAX
Manage the background-tasks component within the security datasync module using the syntax shown in the following sections.

MODIFY
modify background-tasks [task_name]
options:
daily-start-time [time]

edit background-tasks [task_name]
options:
all-properties
non-default-properties

DISPLAY
list background-tasks
list background-tasks [task_name]
show running-config background-tasks
show running-config background-tasks [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line

show background-tasks
show background-tasks [[[name] | [glob] | [regex]] ...]
options:
field-fmt

DESCRIPTION
You can use the background-tasks component to modify or display the Schedule of the Datasync Background Tasks. This controls the Daily Start Time of each task.

The objects cannot be created or deleted, but only modified and displayed. All of the times displayed and configured are in local time. However, they are stored as UTC, so changing the timezone of the BIG-IP will change the display of the configured time.

There is a limitation on the time which can be configured: the new configuration cannot be within 2 hours of the active time. The configured daily-start-time must be before the value of daily-start-time-must-be-before, OR after the value of daily-start-time-must-be-after, inclusive.

These objects are synchronized via the datasync-global-dg device-group.

Use the show command to display the current task schedule. This shows the start time of each task.

OPTIONS

daily-start-time

Specifies the time of day in which the generation will start.

daily-start-time-must-be-before

Displays the upper limitation: daily-start-time must be before this displayed time.

daily-start-time-must-be-after

Displays the lower limitation: daily-start-time must be after this displayed time.

error-msg

Displays the configuration error message which causes daily-start-time to be non-configurable, or "none" if there is no error.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-08-30 security datasync background-tasks(1)

security datasync device-stats

NAME

device-stats - Display the Datasync Framework device stats.

MODULE

security datasync

SYNTAX

Display the device-stats component within the security datasync module using the syntax shown in the following sections.

DISPLAY

list device-stats

list device-stats [[[name] | [glob] | [regex]] ...]

show running-config device-stats

show running-config device-stats [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

recursive

DESCRIPTION

You can use the device-stats component to display the datasync device-stats for that are updated by the Datasync Framework.

OPTIONS

device

Displays the device name reflected by the stats.

table

Displays the table type reflected by the stats.

profile-chksum

Displays the checksum of the current active profile.

activation-epoch

Displays the activation epoch of the current active generation.

rows-available

Displays the number of available rows in the current active buffer.

rows-in-use

Displays the number of rows that are in use in the current active buffer.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh, trust-domain

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-07-10 security datasync device-stats(1)

security datasync global-profile

NAME

global-profile - Manage the Datasync Framework settings that are globally synced across devices.

MODULE

security datasync

SYNTAX

Manage the global-profile component within the security datasync module using the syntax shown in the following sections.

CREATE/MODIFY

create global-profile [name]

modify global-profile [name]

options:

table [name]

activation-epoch [auto | [integer]]

deactivation-epoch [deactivated | always-active | [integer]]

min-rows [default | [integer]]

max-rows [default | [integer]]

regen-time-offset [default | [integer]]

regen-interval [default | none | [integer]]

grace-time [default | [integer]]

master-key [string]

scramble-alg [string]

hash-alg [string]

mac-alg [string]

mode-of-op [string]

rsa-exp [none | rsa-3 | rsa-f4 | default]

rsa-bits [default | none | [integer]]

params [string]

edit global-profile [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list global-profile

list global-profile [[[name] | [glob] | [regex]] ...]

show running-config global-profile

show running-config global-profile [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

recursive

DELETE

delete global-profile [name]

DESCRIPTION

You can use the global-profile component to create, modify, display, or delete a datasync global-profile for use with the Datasync Framework. All global-profiles must reside in the automatically created /Common/datasync-global folder. These profiles are synced across all of the devices in the trust domain using the manually-synced device-group datasync-global-dg.

Warning: Creating, modifying or deleting global-profiles may result in the system being offline, in the case

of a bad configuration.

EXAMPLES

```
create global-profile /Common/datasync-global/my_global_profile table table_name
```

Creates a custom datasync global-profile named my_global_profile for table table_name with initial settings, a random master-key, and an activation-epoch 30 minutes in the future.

```
list global-profile /Common/datasync-global/*
```

Displays the properties of all datasync global-profiles.

OPTIONS

table

Specifies the table to which the profile belongs.

activation-epoch

Specifies the epoch at which the profile becomes active, in UNIX-time.

deactivation-epoch

Specifies the epoch at which the profile becomes inactive, in UNIX-time.

create-timestamp

Displays the timestamp at which the profile was created, in UNIX-time.

min-rows

Specifies the minimum number of rows to generate before going online.

max-rows

Specifies the maximum number of rows to generate.

regen-time-offset

Specifies the time offset at which regeneration will be done, in seconds.

regen-interval

Specifies the time interval at which regeneration will be done, in seconds.

grace-time

Specifies the grace time during which new buffers are supported, but not yet activated, in seconds.

master-key

Specifies the secured master key upon which all cryptography is based. Use \"auto\" to generate a random key.

scramble-alg

Specifies the scrambling algorithm to use.

hash-alg

Specifies the hashing algorithm to use.

mac-alg

Specifies the MAC algorithm to use.

mode-of-op

Specifies the scrambling mode-of-operation to use.

rsa-exp

Specifies the RSA exponent.

rsa-bits

Specifies the number of bits to use for RSA keys.

params

Specifies additional internal parameters.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh, trust-domain

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-07-10 security datasync global-profile(1)

NAME

local-profile - Manage the Datasync Framework settings that are local, and not synced across devices.

MODULE

security datasync

SYNTAX

Manage the local-profile component within the security datasync module using the syntax shown in the following sections.

MODIFY

modify local-profile [table_name]

options:

buf-size [integer]
ds-area [none | asm | fps]
rows-bulk [integer]
gen-timeout-sec [integer]
min-mem-mb [integer]
min-cpu-percent [integer]
max-gen-rows [infinite | [integer]]
keep-conf-files [integer]
gen-pause-sec [integer]
offline-until-gen [enable | disable]
max-iowait-percent [integer]

edit local-profile [table_name]

options:

all-properties
non-default-properties

DISPLAY

list local-profile

list local-profile [table_name]

show running-config local-profile

show running-config local-profile [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition
recursive

DESCRIPTION

You can use the local-profile component to modify or display a datasync local-profile for use with the Datasync Framework. Each profile is for a single table. The profiles cannot be created or deleted, only modified. These profiles are local; they are not synced across devices.

OPTIONS

buf-size

Specifies the size of each buffer in the table, in bytes.

ds-area

Specifies the memory area on which the table will be allocated.

rows-bulk

Specifies the number of rows to generate in each bulk.

gen-timeout-sec

Specifies the timeout of running the external generator per single bulk.

min-mem-mb

Specifies the minimum available memory in MB to start generator in non-urgent mode.

min-cpu-percent

Specifies the minimum available CPU percent to start generator in non-urgent mode.

max-gen-rows

Specifies the maximum rows to generate.

keep-conf-files

Specifies the number of configuration files to keep when rolling old ones.

gen-pause-sec

Specifies the time in seconds to pause between each bulk generation when in non-urgent mode.

offline-until-gen

Specifies if the system should be offline until the table is generated.

min-cpu-percent

Specifies the maximum iowait percent to start generator in non-urgent mode.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2019-07-01 security datasync local-profile(1)

security debug drop-redirect-stats

NAME

drop-redirect-stats - Configures Debuggability drop redirect mode.

MODULE

security firewall

SYNTAX

Configure drop redirect feature or display stats using the following syntax.

DISPLAY

show drop-redirect-stats

MODIFY

reset-stats drop-redirect-stats

DESCRIPTION

This command displays size and number of packets that were drop redirected.

EXAMPLES

show security debug drop-redirect-stats

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2017. All rights reserved.

BIG-IP 2017-07-20 security debug drop-redirect-stats(1)

security debug matcher

NAME

debug - Configures Debuggability drop redirect mode.

MODULE

security firewall

SYNTAX

Configure drop redirect feature or display stats using the following syntax.

MODIFY

```
modify debug
matcher {
  drop-redirect {
    drop-redirect-mode {
      disable
      redirect-all
      redirect-hw-only
      redirect-sw-only
    }
  }
}
```

DISPLAY

```
show debug
drop-redirect-stats
```

DESCRIPTION

Debuggability drop redirection feature redirects HW dropped packets to a specified interface. This interface may be set using sys db variable debug.hwdropredirect.interface. The feature can also redirect only certain types of drops. This can be done by using sys db variable debug.doshwdropredirect.disable.

Full List of HW Redirect Modes # Disable GlobalDoSVector drop redirects bit-0 # Disable sPVADoSVector

drop redirects bit-1 # Disable sPVAIPBlacklist drop redirects bit-2 # Disable sPVAIPRateLimit drop
redirects bit-3 # Disable NeuronBlacklist drop redirects bit-4 # Disable DuplicateSYN drop redirects
bit-5

Once an interface is set-up, redirect-hw-only mode can be enabled as the following example.

EXAMPLES

```
modify security debug matcher drop-redirect drop-redirect-mode redirect-hw-only
```

Configures dropped packets to be redirected to a specified interface.

```
BIG-IP 2018-01-10 security debug matcher(1)
```

security debug register

NAME

register - Configures a debug register.

MODULE

security debug

SYNTAX

Configure the register component within the security debug module using the syntax in the following sections.

MODIFY

```
modify register [name]
```

options:

```
all  
description [string]  
destination {  
  address [ip_address/prefixlen]  
  port [port]  
}  
[disabled | enabled]  
match-ip-version [false | true]  
protocol [any | [protocol] ]  
source {  
  address [ip_address/prefixlen]  
  port [port]  
  vlan [vlan name]  
}
```

```
edit register [ [name] ... ]
```

options:

```
all-properties  
non-default-properties
```

```
reset-stats register [ [name] ... ]
```

DISPLAY

```
list register
```

```
list register [ [name] ... ]
```

options:

```
all-properties  
non-default-properties  
one-line
```

```
show register [ [name] ... ]
```

options:

```
all-properties (default | exa | gig | kil | meg | peta | raw | tera |  
  yotta | zetta)  
field-fmt
```

RUN

```
run register [name]
```

options:

```
filename [filename | stdout]  
max-file-mb [integer]  
max-packets [integer]  
unidirectional [true | false]  
capture-start  
capture-stop
```

DESCRIPTION

This component configures the traffic flow for hardware debug functionality based on the incoming packets' IP header 6-tuple values.

The run command performs the hardware debug functionality by capturing the network traffic which matches the register configuration.

EXAMPLES

```
modify register r1 enabled source { address 1.1.1.0/24 port any vlan vlan-168 }
```

Configure register r1 to match the traffic from address 1.1.1.0/24 on any port and vlan defined in vlan-168.

```
reset-stats register r1
```

Reset the statistics of register named r1.

```
show register r1
```

Displays statistics and status of register named r1.

```
run register r1 filename /shared/f1.cap max-file-mb 100 max-packets 5000 capture-start
```

Start capturing the network traffic based on r1 configuration and save the tcpdump file to /shared/f1.cap, up to 100 mb or 5000 packets whichever comes first.

```
run register r1 capture-stop
```

Stop capturing the network traffic on r1 manually.

OPTIONS

description

User-defined description.

destination

Matches against each packet's destination IP and/or destination port.

address

Specifies an IP address and network to compare against the packet's destination address.

The format for an IPv4 address is a.b.c.d[/prefix]. The general format for an IPv6 address is a:b:c:d:e:f:g:h[/prefix]; you can shorten this by eliminating leading zeros from each field (for example, you can shorten "2001:0db7:3f4a:09dd:ca90:ff00:0042:8329" to "2001:db7:3f4a:9dd:ca90:ff00:42:8329"), and/or by removing the longest contiguous field of zeros (for example, you can shorten "2001:0:0:0:c34a:0:23ff:678" to "2001::c34a:0:23ff:678"). TMSH accepts any valid text representation of IPv6 addresses, as defined in RFC 2373 (see).

port Specifies a port to compare against the packet's destination port.

match-ip-version

Specifies whether any/any6 in source/destination address is to match 'any IPv4', 'any IPv6', or 'any IPv4 and any IPv6' addresses.

If match-ip-version is true, both source and destination addresses must have the same IP address family. If match-ip-version is false and both source and destination addresses are any or any6, both addresses represent 'any IPv4 and IPv6 addresses'.

If match-ip-version is false and only one address is set to any or any6, the address is interpreted based on the other IP address' family (IPv4 or IPv6). The default is false.

protocol

Specifies the IP protocol to compare against the packet. The default value is any.

source

Matches against each packet's source IP, source port, and/or source VLAN.

address

Specifies an IP address and network to compare against the packet's source address.

The format for an IPv4 address is a.b.c.d[/prefix]. The general format for an IPv6 address is a:b:c:d:e:f:g:h[/prefix]

port Specifies a port to compare against the packet's source port.

vlan Specifies a vlan name.

filename

Specifies the full path of the file in which to capture the packets from the run command.

The option is only for the run command with option capture-start. The default value is stdout if it's not specified. The tcpdump will be displayed on the console if the value is stdout.

max-file-mb

Specifies the maximum file size in the run command.

The option is only for the run command with option capture-start. The default value is 1 if it's not specified. The unit is in 1,000,000 bytes.

max-packets

Specifies the maximum number of packets that can be captured in the run command.

The option is only for the run command with option capture-start. The default value is 1000 if it's not specified.

unidirectional

Specifies that only the unidirectional traffic can be captured in the run command.

The option is only for the run command with option capture-start. The default value is false which means bidirectional traffic will be captured. If the value is true, the command will only capture unidirectional traffic.

capture-start

Specifies the action to start capturing the network traffic.

The option is only for the run command. Either capture-start or capture-stop must be specified as the last option on the run command.

capture-stop

Specifies the action to stop capturing the network traffic.

The option is only for the run command. Either capture-start or capture-stop must be specified as the last option on the run command.

SEE ALSO

edit, list, modify, security, debug, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-12-05 security debug register(1)

security device-id attribute

NAME

device-id - Configures Device-ID collected attributes.

MODULE

security

SYNTAX

Configure the device-id component within the security module using the syntax shown in the following sections.

MODIFY

```
modify attribute [name]
options:
  collect {
[enabled | disabled]
  }
```

DISPLAY

```
list device-id
```

DESCRIPTION

You can use the device-id component to modify or display a Device-ID attribute collection.

EXAMPLES

```
modify device-id attribute att05 {collect disabled}
```

Disables collection of att05 by Device-ID JavaScript.

```
list device-id
```

Displays the properties of all device-id attributes.

OPTIONS

attribute

Specifies the Device-ID category to be configured.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2016. All rights reserved.

BIG-IP 2017-05-24 security device-id attribute(1)

security device device-context

NAME

device-context - Configures the security device context policies.

MODULE

security device

SYNTAX

MODIFY

```
modify device-context
options:
description [string]
nat-policy [ [policy_name] | none ]
```

```
edit device-context
```

options:

```
all-properties
non-default-properties
```

```
reset-stats device-context
```

```
nat-rules { [rule name] }
```

DISPLAY

```
list device-context
```

```
show running-config device-context nat-rules
```

```
show device-context nat-rules
```

DESCRIPTION

You can use the device-context component to configure device level (a.k.a global) policies.

EXAMPLES

```
modify device-context nat-policy policy1
```

Configures nat policy named policy1 to be used as the device level nat policy.

```
list device-context
```

```
security device device-context {
nat-policy /Common/policy1
}
```

Displays the current list of policies at device context level.

OPTIONS

description

Specifies nat policies on device context level.

nat-policy

Specifies a NAT policy. This is the policy used to match incoming traffic to a virtual server and perform source/destination translations (as per the matched NAT rule) if and only if: virtual server does not have a NAT policy configured (see ltm virtual)

AND virtual server has nat-policy.use-route-domain-policy disabled

OR virtual server has nat-policy.use-route-domain-policy enabled but the route domain does not have a NAT policy configured

AND virtual server has nat-policy.use-device-policy enabled.

nat-rules

Specifies security nat rules enforced on device context level via referenced nat-policy.

SEE ALSO

edit, list, modify, security firewall address-list, security firewall port-list, security nat policy, security log profile, security nat source-translation, security nat destination-translation, ltm virtual, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015. All rights reserved.

BIG-IP 2016-03-14 security device device-context(1)

security dos auto-thresholds heavy-urls

NAME

heavy-urls - Displays the heavy URL auto-calculated thresholds.

MODULE

security dos auto-thresholds

SYNTAX

Display information about the heavy-urls component within the security dos auto-thresholds module using the following syntax.

DISPLAY

show heavy-urls

options:

all
field-fmt
recursive

DESCRIPTION

You can use the heavy-urls component to display the heavy URL auto-calculated thresholds.

EXAMPLES

```
show heavy-urls my_dos_profile
```

Display the heavy URL auto-calculated thresholds for a DoS profile named my_dos_profile.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-06-22 security dos auto-thresholds heavy-urls(1)

security dos auto-thresholds stress-based

NAME

stress-based - Displays the Stress-based auto-calculated thresholds.

MODULE

security dos auto-thresholds

SYNTAX

Display information about the stress-based component within the security dos auto-thresholds module using the following syntax.

DISPLAY

show stress-based

options:

all
field-fmt
recursive

DESCRIPTION

You can use the stress-based component to display the Stress-based auto-calculated thresholds.

EXAMPLES

```
show stress-based my_dos_profile
```

Display the Stress-based auto-calculated thresholds for a DoS profile named my_dos_profile.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

security dos auto-thresholds top-device-ids

NAME

top-device-ids - Displays the top device ID auto-calculated thresholds.

MODULE

security dos auto-thresholds

SYNTAX

Display information about the top-device-ids component within the security dos auto-thresholds module using the following syntax.

DISPLAY

show top-device-ids

options:

all

field-fmt

recursive

DESCRIPTION

You can use the top-device-ids component to display the top device ID auto-calculated thresholds.

EXAMPLES

show top-device-ids my_dos_profile

Display the top device ID auto-calculated thresholds for a DoS profile named my_dos_profile.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-06-22 security dos auto-thresholds top-device-ids(1)

security dos auto-thresholds top-geolocations

NAME

top-geolocations - Displays the top geolocation auto-calculated thresholds.

MODULE

security dos auto-thresholds

SYNTAX

Display information about the top-geolocations component within the security dos auto-thresholds module using the following syntax.

DISPLAY

show top-geolocations

options:

all

field-fmt

recursive

DESCRIPTION

You can use the top-geolocations component to display the top geolocation auto-calculated thresholds.

EXAMPLES

show top-geolocations my_dos_profile

Display the top geolocation auto-calculated thresholds for a DoS profile named my_dos_profile.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-06-22 security dos auto-thresholds top-geolocations(1)

security dos auto-thresholds top-source-ips

NAME

top-source-ips - Displays the top source IP auto-calculated thresholds.

MODULE

security dos auto-thresholds

SYNTAX

Display information about the top-source-ips component within the security dos auto-thresholds module using the following syntax.

DISPLAY

show top-source-ips

options:

all

field-fmt

recursive

DESCRIPTION

You can use the top-source-ips component to display the top source IP auto-calculated thresholds.

EXAMPLES

```
show top-source-ips my_dos_profile
```

Display the top source IP auto-calculated thresholds for a DoS profile named my_dos_profile.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-06-22 security dos auto-thresholds top-source-ips(1)

security dos auto-thresholds top-urls

NAME

top-urls - Displays the top URL auto-calculated thresholds.

MODULE

security dos auto-thresholds

SYNTAX

Display information about the top-urls component within the security dos auto-thresholds module using the following syntax.

DISPLAY

show top-urls

options:

all
field-fmt
recursive

DESCRIPTION

You can use the top-urls component to display the top URL auto-calculated thresholds.

EXAMPLES

```
show top-urls my_dos_profile
```

Display the top URL auto-calculated thresholds for a DoS profile named my_dos_profile.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-06-22 security dos auto-thresholds top-urls(1)

security dos auto-thresholds tps-based

NAME

tps-based - Displays the TPS-based auto-calculated thresholds.

MODULE

security dos auto-thresholds

SYNTAX

Display information about the tps-based component within the security dos auto-thresholds module using the following syntax.

DISPLAY

```
show tps-based
```

options:

all
field-fmt
recursive

DESCRIPTION

You can use the tps-based component to display the TPS-based auto-calculated thresholds.

EXAMPLES

```
show tps-based my_dos_profile
```

Display the TPS-based auto-calculated thresholds for a DoS profile named my_dos_profile.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-06-22 security dos auto-thresholds tps-based(1)

security dos autodos-file-object

NAME

autodos-file-object - autodos internal file object.

MODULE
security

SYNTAX
Configure autodos-file-object within security dos module using the syntax shown in the following sections.

CREATE/MODIFY
create autodos-file-object [name]
modify autodos-file-object [name]
options:
app-service [[string] | none]
source-path [file name]

DELETE
delete autodos-file-object [name]

DESCRIPTION
autodos file object is used internally for synchronization of dos module data across HA (high availability) or Cluster setup.

EXAMPLES
create name source-path /var/dos/file.json

Creates a file object named name, that gets its contents from source file /var/dos/file.json.

delete filename

Deletes the file object named filename.

OPTIONS
app-service
Specifies the name of the application service to which the file object belongs. The default value is none.

source-path
The name of the source file from which the contents are copied to create the file object.

SEE ALSO
create, delete, edit, modify, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2017-07-07 security dos autodos-file-object(1)

security dos behavioral-signature

NAME
behavioral-signature - Configures a Behavioral Bot Signature.

This component has been deprecated and replaced by dos-signature in 13.1.0.

MODULE
security dos

SYNTAX
Configure the behavioral-signature component within the security dos module using the syntax shown in the following sections.

MODIFY
modify behavioral-signature [name]
options:
alias [name]
status [approved | new | revoked }

DISPLAY
list behavioral-signature

DELETE
delete behavioral-signature [name]

DESCRIPTION
You can use the behavioral-signature component to modify, display, or delete a Behavioral Signature.

EXAMPLES

list behavioral-signature

Displays the properties of all Behavioral Signatures.

OPTIONS

alias

Specifies the behavioral signature user defined alias.

status

Specifies the behavioral signature status.

SEE ALSO

edit, list, modify, security, security dos, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-09 security dos behavioral-signature(1)

security dos bot-signature-category

NAME

bot-signature-category - Configures the Bot Signature Categories.

MODULE

security dos

SYNTAX

Configure the bot-signature-category component within the security dos module using the syntax shown in the following sections.

CREATE/MODIFY

create bot-signature-category [string]

modify bot-signature-category [name]

options:

type [benign | malicious]

DISPLAY

list bot-signature-category

DELETE

delete bot-signature-category [name]

DESCRIPTION

You can use the bot-signature-category component to create, modify, display, or delete a Bot Signature Category.

EXAMPLES

create bot-signature my_signature_category

Creates a custom Bot Signature Category named my_signature_category with initial settings.

list bot-signature-category

Displays the properties of all Bot Signature Categories.

OPTIONS

type Specifies the bot signature type. The possible values are benign or malicious.

SEE ALSO

edit, list, modify, security, security dos, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-10-22 security dos bot-signature-category(1)

security dos bot-signature

NAME

bot-signature - Configures the Bot Signatures.

MODULE

security dos

SYNTAX

Configure the bot-signature component within the security dos module using the syntax shown in the following sections.

CREATE/MODIFY

```
create bot-signature [string]
modify bot-signature [name]
options:
  category [name]
  domains [none | add | delete | replace-all-with] { [string] ... }
  risk [high | low | medium]
  rule [string]
  signature-references [string]
  url {
  match-type [contains | regexp]
  search-string [string]
  }
  user-agent {
  match-type [contains | regexp]
  search-string [string]
  }
```

DISPLAY

```
list bot-signature
```

DELETE

```
delete bot-signature [name]
```

DESCRIPTION

You can use the bot-signature component to create, modify, display, or delete a Bot Signature.

EXAMPLES

```
create bot-signature my_signature
```

Creates a custom Bot Signature named my_signature with initial settings.

```
list bot-signature
```

Displays the properties of all Bot Signatures.

OPTIONS

category

Specifies the bot signature category.

domains

Specifies the bot signature domain names.

risk

Specifies the bot signature risk. The possible values are high, low and medium.

rule

Specifies the bot signature rule.

signature-references

Specifies the bot signature references.

url

Specifies the bot signature's url matching rule. The following options are available:

match-type

Specifies the bot signature's url rule match type. The possible values are contains or regexp.

search-string

Specifies the bot signature's url string that should be matched.

user-agent

Specifies the bot signature's user-agent matching rule. The following options are available:

match-type

Specifies the bot signature's user-agent rule match type. The possible values are contains or regexp.

search-string

Specifies the bot signature's user-agent string that should be matched.

SEE ALSO

edit, list, modify, security, security dos, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-11-27 security dos bot-signature(1)

security dos device-config

NAME

device-config - Configures the global network DoS profile.

MODULE

security dos

SYNTAX

Configure the global network DoS profile component within the security dos module using the syntax shown in the following sections.

MODIFY

modify device-config dos-device-config

options:

auto-threshold-sensitivity [field deprecated since 13.0.0]

ip-uncommon-protolist [string]

threshold-sensitivity [low | medium | high]

custom-signatures [none | add | delete | modify | replace-all-with] {

name [string] {

options:

manual-detection-threshold [integer]

manual-mitigation-threshold [integer]

state [disabled | learn-only | detect-only | mitigate]

threshold-mode [fully-automatic | manual | manual-multiplier-mitigation | stress-based-mitigation]

}

}

dos-device-vector {

[vector type] {

allow-advertisement [disabled | enabled]

allow-upstream-scrubbing [disabled | enabled]

attacked-dst [disabled | enabled]

auto-blacklisting [enabled | disabled]

auto-scrubbing [disabled | enabled]

auto-threshold [disabled | enabled]

bad-actor [disabled | enabled]

blacklist-category [enter name of ip-intelligence category]

blacklist-detection-seconds [integer]

blacklist-duration [integer]

ceiling [integer | infinite]

default-internal-rate-limit [integer | infinite]

detection-threshold-percent [integer | infinite]

detection-threshold-pps [integer | infinite]

enforce [enabled | disabled] [field deprecated since 13.1.0]

floor [integer]

multiplier_mitigation_percentage [integer]

packet-types [add | delete | replace-all-with] {

[atomic-frag | bad-packet | dns-a-query | dns-a-query | dns-aaaa-query |

dns-any-query | dns-axfr-query | dns-cname-query | dns-ixfr-query |

dns-mx-query | dns-ns-query | dns-other-query | dns-oversize |

dns-ptr-query | dns-response-flood | dns-soa-query | dns-srv-query |

dns-txt-query | exthdr | host-unrch | igmp | ip-overlap-frag |

ipfrag | ipv4-all | ipv4-any-other | ipv4-icmp | ipv6-all |

ipv6-any-other | ipv6-icmp | no-l4 | rthdr0 | sip-ack-method |

sip-bye-method | sip-cancel-method | sip-invite-method |

sip-malformed | sip-message-method | sip-notify-method |

sip-options-method | sip-other-method | sip-prack-method |

sip-publish-method | sip-register-method | sip-subscribe-method | sip-uri-limit |

suspicious | tcp-bad-ack | tcp-psh-flood | tcp-rst | tcp-syn-only |

tcp-synack | tcp-winsize | tidcmp | udp]

packet-types none

per-dst-ip-detection-pps [integer]

per-dst-ip-limit-pps [integer]

per-source-ip-detection-pps [integer]

per-source-ip-limit-pps [integer]

scrubbing-category [enter name of scrubbing category | "none"]

scrubbing-detection-seconds [integer]

scrubbing-duration [integer]

simulate-auto-threshold [enable | disable]

```

state [disabled | learn-only | detect-only | mitigate]
suspicious [ false | true ]
threshold-mode [manual | stress-based-mitigation | fully-automatic | manual-multiplier-mitigation]
valid-domains [add | delete | replace-all-with] {
  [domain names] ...
}
valid-domains none
}
}
dynamic-signatures {
detection [disabled | enabled | learn-only] [field deprecated since 13.1.0]
mitigation [none | low | medium | high] [field deprecated since 13.1.0]
scrubber-advertisement-period [integer] [field deprecated since 13.1.0]
scrubber-category [name] [field deprecated since 13.1.0]
scrubber-enable [yes | no] [field deprecated since 13.1.0]
network {
  detection [disabled | enabled | learn-only]
  mitigation [none | low | medium | high | manual-multiplier]
  scrubber-advertisement-period [integer]
  scrubber-category [name]
  scrubber-enable [yes | no]
}
}
dns {
  detection [disabled | enabled | learn-only]
  mitigation [none | low | medium | high | manual-multiplier]
}
}
dns-dos-mitigation-percentage [integer]
log-publisher [name]
network-dos-mitigation-percentage [integer]
sip-dos-mitigation-percentage [integer]
syn-cookie-dsr-flow-reset-by [bigip | client | none]
syn-cookie-whitelist [disabled | enabled]
tscookie-vlans
[add | delete | replace-all-with] {
[vlan name] ...
}
tscookie-vlans [default | none]

```

```

reset-stats device-config dos-device-config
options:
  dns-nxdomain-stat

```

```

DISPLAY
list device-config dos-device-config
show running-config device-config dos-device-config
options:
  all-properties
  non-default-properties
  one-line

```

```

show device-config dos-device-config
options:
  dns-nxdomain-stat
  field-fmt
  query-valid-domain [domain-name]

```

```

RUN
run device-config
options:
  auto-threshold-relearn
  dns-nxdomain-relearn
  dynamic-signatures-history-relearn

```

DESCRIPTION

This component is used to modify or display the global device DoS profile and statistics for use with network DoS Protection functionality.

EXAMPLES

```
modify device-config ...
```

Modifies the global DoS profile settings.

```
list device-config
```

Displays all the properties of the device DoS profile.

```
run device-config dos-device-config auto-threshold-relearn
```

Clears the auto-threshold history for all the device auto-threshold vectors.

```
run device-config dos-device-config dns-nxdomain-relearn
```

Clears the dns-nxdomain history for all the device dns-nxdomain vectors.

```
run device-config dos-device-config dynamic-signatures-history-relearn
```

Clears the dynamic-signatures history for all the device dynamic-signatures vectors.

show device-config dos-device-config dns-nxdomain-stat

Displays the dns-nxdomain statistics for the device.

reset-stats device-config dos-device-config dns-nxdomain-stat

Resets the dns-nxdomain statistics for the device.

OPTIONS

auto-threshold-sensitivity

This option is deprecated in version 13.0.0.

dos-device-vector

Configures attack detection thresholds and rate limit parameters for network DoS vectors.

log-publisher

Specifies the name of the log publisher which logs translation events. See help sys log-config for more details on the logging sub-system.

ip-uncommon-protolist

Specifies the name of an IP uncommon protocol list component. The default is /Common/ip-uncommon-protolist. This is ready-only field.

threshold-sensitivity

Specifies the guidance on how aggressively (how much to pad) to adjust the "Detection/Rate-limit Threshold". Available settings are low, medium and high. This setting is used for Autodos and Behavioral DoS features. Default is set to medium.

network-dos-mitigation-percentage

Specifies the mitigation multiplier value of all the device network dos vector in percentage in the manual-multiplier-mitigation mode.

dns-dos-mitigation-percentage

Specifies the mitigation multiplier value of all the device dns dos vector in percentage in the manual-multiplier-mitigation mode.

sip-dos-mitigation-percentage

Specifies the mitigation multiplier value of all the device sip dos vector in percentage in the manual-multiplier-mitigation mode.

syn-cookie-dsr-flow-reset-by

Specifies how TCP SYN Flood is handled when syn-cookie-whitelist is enabled and the attack is detected in Direct Server Return(DSR) mode. The default value is none.

syn-cookie-whitelist

Specifies whether or not to use a SYN Cookie WhiteList when doing software SYN Cookies. This means not doing a SYN Cookie for the same src IP address if it has been done already in the previous tm.flowstate.timeout (30) seconds. The default value is disabled.

dynamic-signatures

Specifies options related to L4-L7 Behavioral DoS (Dynamic Signatures) feature that is applicable at the global/device level. These settings are used to learn the characteristic of the traffic at the device level (across all domains and virtual servers) and generate dynamic signatures as applicable to detect and mitigate anomalous traffic.

Following options are configurable for this feature at global/device level:

network

detection

Specifies the mode for detection of anomalies in traffic for the purpose of dynamic signature generation. Following modes are supported: disabled, enabled and learn-only.

Mode learn-only is same as enabled except that the system does not generate any logs (or alerts the user). It is used mainly to learn the baseline thresholds for the traffic.

Default is disabled.

mitigation

Specifies the mode for mitigation of anomalous traffic (specified in form of dynamic signatures). Following modes are supported: none, low, medium and high.

Each mode represents the severity (or aggressiveness) at which the system should try to mitigate the anomalous traffic.

Default is none.

multiplier-mitigation-percentage

Specifies the mitigation multiplier value of this specific dos signature in percentage when using manual-multiplier-mitigation mode. The default value is inherited from the corresponding device level/profile mitigation multiplier value of the same dos family.

scrubber-enable

Specifies the configuration mode for enabling or disabling the feature to scrub the attack traffic upon dynamic signature match. Default is no.

scrubber-category

Specifies the IP Intelligence category used for scrubbing the attack traffic upon dynamic signature match that constitutes destination IP address component. Default category is `attacked_ips`.

scrubber-advertisement-period

Specifies the advertisement period for which the attack traffic is scrubbed. Default is 300 seconds.

dns detection

Specifies the mode for detection of anomalies in traffic for the purpose of dynamic signature generation. Following modes are supported: `disabled`, `enabled` and `learn-only`.

Mode `learn-only` is same as `enabled` except that the system does not generate any logs (or alerts the user). It is used mainly to learn the baseline thresholds for the traffic.

Default is `disabled`.

mitigation

Specifies the mode for mitigation of anomalous traffic (specified in form of dynamic signatures). Following modes are supported: `none`, `low`, `medium` and `high`.

Each mode represents the severity (or aggressiveness) at which the system should try to mitigate the anomalous traffic.

Default is `none`.

custom-signatures

Specifies options related to L4 Behavioral DoS Signatures feature that is applicable at the global/device level. Signatures can be added to a `dos-profile` and the signature criteria will be used for detection and mitigation of anomalous traffic.

Following options are configurable for each signature added:

threshold-mode

Specifies the mode for setting the rate limit thresholds to be used for the matching traffic. Following modes are supported: `manual`, `fully-automatic`, `manual-multiplier-mitigation` and `stress-based-mitigation`. Default is `manual`.

state

Specifies the operational state of the attached signature. The states supported are: `disabled`, `learn-only`, `detect-only` and `mitigate`. Default is `disabled`.

manual-detection-threshold

Specifies manual detection threshold for a custom signature. It is applicable only if `threshold-mode` is set to either `manual` or `stress-based-mitigation`

Default is infinite.

manual-mitigation-threshold

Specifies manual mitigation threshold for a custom signature. It is applicable only if `threshold-mode` is set to either `manual` or `stress-based-mitigation`

Default is infinite.

tscookie-vlans

Specifies the VLANs on which we will do TCP timestamp cookie based validation of TCP ACK packets and use the TCP BAD ACK DoS vector to mitigate a TCP ACK flood attack.

VECTOR TYPES

arp-flood

ARP Flood.

bad-ext-hdr-order

IPv6 extension headers in packet are out of order.

bad-icmp-chksum

Bad ICMP checksum.

bad-icmp-frame

Bad ICMP frames. To see the various reasons why ICMP frames are classified as bad, please refer to the written documentation.

bad-igmp-frame

Bad IGMP frames. To see the various reasons why IGMP frames are classified as bad, please refer to the written documentation.

bad-ip-opt

IPv4 option with illegal length.

bad-ipv6-hop-cnt

Bad IPv6 hop count. Terminated packet (`cnt==0`). Dropped when the rate hits rate limit.

bad-ipv6-ver

Bad IPv6 version. IP Version in the IPv6 packet is not 6.

bad-sctp-chksum
Bad SCTP Checksum type.

bad-tcp-chksum
Bad TCP checksum.

bad-tcp-flags-all-clr
Bad TCP flags (all TCP header flags cleared).

bad-tcp-flags-all-set
Bad TCP flags (all flags set).

bad-ttl-val
Bad IP TTL value (TTL == 0 for IPv4).

bad-udp-chksum
Bad UDP checksum.

bad-udp-hdr
Bad UDP header. To see the various reasons why UDP headers are classified as bad, please refer to the written documentation.

bad-ver
Bad IP version 4. IPv4 version in IP header is not 4.

dns-a-query
DNS A query packet.

dns-aaaa-query
DNS AAAA query packet.

dns-any-query
DNS any query packet.

dns-axfr-query
DNS AXFR query packet.

dns-cname-query
DNS CNAME query packet.

dns-ixfr-query
DNS IXFR query packet.

dns-malformed
DNS Malformed packet.

dns-mx-query
DNS MX query packet.

dns-ns-query
DNS NS query packet.

dns-nxdomain-query
DNS NXDOMAIN query packet.

dns-other-query
DNS OTHER query packet.

dns-oversize
DNS packet with size > . This sys db tunable is configurable with Dos.MaxDNSframeSize.

dns-ptr-query
DNS PTR query packet.

dns-qdcount-limit
DNS QDCOUNT LIMIT query packet.

dns-response-flood
DNS RESPONSE FLOOD query packet.

dns-soa-query
DNS SOA query packet.

dns-txt-query
DNS TXT query packet.

dns-srv-query
DNS SRV query packet.

dup-ext-hdr
Duplicate IPv6 extension headers.

ether-brdcst-pkt
Ethernet broadcast packet.

ether-mac-sa-eq-da
Ethernet MAC SA == DA.

ether-multicast-pkt
Ethernet multicast packet.

ext-hdr-too-large
IPv6 extension header size too large. The max IPV6 extension header size is configurable via the sys db variable dos.maxipv6extsize.

fin-only-set
TCP header with only the FIN flag set.

flood
A Flood is an attack where multiple (typically many) endpoints initiate network traffic to a single subnet or receiving endpoint.

hdr-len-gt-l2-len
Header length > L2 length. No room in L2 packet for IPv4 header (including options).

hdr-len-too-short
Header length too short. IPv4 header length in IP header is less than 20 bytes.

hop-cnt-leq-one
IPv6 hop count <= and the packet needs to be forwarded. This sys db tunable is configurable by the sys db variable tm.minipv6hopcnt.

host-unreachable
ICMP packets of type "Host Unreachable".

icmp-frag-flood
ICMP fragments flood.

icmp-frame-too-large
Packets larger than the maximum ICMP frame size. The max ICMP frame size is configurable via the sys db variable dos.maxicmpframesize.

icmpv4-flood
ICMPv4 Flood.

icmpv6-flood
ICMPv6 Flood.

igmp-flood
IGMP Flood.

igmp-frag-flood
IGMP Fragment Flood.

ip-bad-src
IP addr is a broadcast or multicast address.

ip-err-chksum
IP error checksum. IPv4 header checksum error.

ip-frag-flood
IPv4 fragment flood.

ip-len-gt-l2-len
IP length > L2 length. Total length in IPv4 header is greater than the L3 part length in L2 packet.

ip-overlap-frag
IPv4 overlapping fragments.

ip-short-frag
IPv4 fragments whose payload size is less than the minimum IPv4 Fragment size. The minimum size is configurable via the db variable tm.minipfragsize.

ip-unk-prot
IP Unknown Protocol type.

ip-opt-frames
IP option frames. IPv4 packets with options. db variable tm.acceptipoptions must be enabled to receive IP options.

ip-other-frag
The total IPv4 fragments' size has exceeded the reassembly queue or the maximum IP packet size.

ipv6-atomic-frag
IPv6 frame with frag extension hdr, but the MF and offset fields are both 0.

ipv6-bad-src
IPv6 src address is a multicast address or IPv6 src or destination address is a IPv4 mapped IPv6 address.

ipv6-ext-hdr-frames
IPv6 extended header frames.

ipv6-frag-flood
IPv6 fragment flood.

ipv6-len-gt-l2-len
IPv6 length > L2 length.

ipv6-other-frag
The total IPv6 fragments' size has exceeded the reassembly queue or the maximum IP packet size.

ipv6-overlap-frag
IPv6 overlapping fragments.

ipv6-short-frag
IPv6 fragments whose payload size is less than the minimum IPv6 Fragment size. The minimum size is configurable via the db variable tm.minipv6fragsize.

ipv4-mapped-ipv6
IPv4 mapped IPv6 addresses.

land-attack
Land Attack. IP Src Address equals IP Dst Address. Both V4 and V6 are counted.

l2-len-ggt-ip-len
L2 length >> IP length. L2 packet length is much greater than payload length in IPv4 (L2 length > IP length and L2 length > minimum packet size).

l4-ext-hdrs-go-end
No L4 (extended headers go to or past the end of frame).

no-l4
No L4. No L4 payload for IPv4.

opt-present-with-illegal-len
TCP Option present with illegal length.

payload-len-ls-l2-len
Payload length < L2 length. Payload length in IPv6 header is less than L3 part length in L2 packet.

routing-header-type-0
Routing header type 0 present.

sip-malformed
SIP malformed packet

sip-invite-method
SIP INVITE method packet.

sip-ack-method
SIP ACK method packet.

sip-options-method
SIP OPTIONS method packet.

sip-bye-method
SIP BYE method packet.

sip-cancel-method
SIP CANCEL method packet.

sip-register-method
SIP REGISTER method packet.

sip-publish-method
SIP PUBLISH method packet.

sip-notify-method
SIP NOTIFY method packet.

sip-subscribe-method
SIP SUBSCRIBE method packet.

sip-message-method
SIP MESSAGE method packet.

sip-prack-method
SIP PRACK method packet.

sip-uri-limit
Limit SIP URI length.

sip-other-method
SIP OTHER method packet.

sweep
A Sweep is an attack where a single endpoint initiates network traffic to a large number of receiving endpoints or subnets.

syn-and-fin-set
SYN && FIN set.

tcp-ack-flood

TCP packets with the ACK flag set (for non-existing flows).

tcp-bad-urg

TCP packets with the URG flag set but URG pointer is 0.

tcp-hdr-len-gt-l2-len

TCP header length > L2 length. No room in packet for TCP header (including options).

tcp-hdr-len-too-short

TCP header length too short (length < 5). The offset field in TCP header is less than 20 bytes.

tcp-opt-overflow-tcp-hdr

TCP option overruns TCP header.

tcp-syn-flood

TCP header with only the SYN flag set.

tcp-synack-flood

TCP header with only the SYN and ACK flags set.

tcp-rst-flood

TCP header with only the RST flag set.

tcp-psh-flood

TCP header with PUSH flag set.

tcp-window-size

TCP non-RST pkt with window size < . This sys db tunable is configurable with Dos.TcpLowWindowSize.

tidcmp

ICMP source quench packets.

too-many-ext-hdrs

Too many extended headers. The IPv6 extended headers are more than 4. This number can be set through db variable dos.maxipv6exthdrs.

tcp-syn-oversize

TCP data-SYN with pktlength > dos.maxsynsize which is 128 bytes by default.

tth-leq-one

TTL <= . For IPv4 forwarding. This sys db tunable is configurable by tm.minipttl.

unk-tcp-opt-type

Unknown TCP option type.

udp-flood

UDP Flood.UDP flood vector counts any UDP packets that either match the UDP Port InclusionList or do not match the UDP Port ExclusionList. "tmsh modify security dos udp-portlist" can be used to configure the udp port list.For more info about udp portlist and how to configure it use "help security dos udp-portlist"

unk-ipopt-type

Unknown IP option type.

ip-uncommon-prot

ip-uncommon-prot vectors counts packets whose protocol is specified in configured ip-uncommon-protolist.

PARAMETERS

allow-advertisement

Enables allow advertisement. The default is disabled.

allow-upstream-scrubbing

Enables allow upstream scrubbing. The default value is disabled.

attacked-dst

Enables attacked-destination. The default value is disabled.

auto-blacklisting

Enables automatic blacklisting of offending source IPs. The default value is disabled.

auto-scrubbing

Enables specifying destination IP scrubbing. The default value is disabled.

auto-threshold

This option is deprecated in version 13.1.0 and is replaced by threshold-mode. Enables the auto threshold mode for dos detection and dos mitigation. The default value is disabled.

bad-actor

Enables per-source IP based bad actor detection. The default value is disabled.

blacklist-category

Blacklist category (of IP intelligence) to which this IP should be added. The default value is none.

blacklist-detection-seconds

Duration in seconds for which the IP has been offending. The default value is 60.

`blacklist-duration`

Duration in seconds for which this IP should be blocked. The default value is 14400.

`ceiling`

Option to set a maximum value ("ceiling") for the default-internal-rate-limit for this vector. The range is from 0 to infinity.

`default-internal-rate-limit`

This parameter is programmed in hardware to limit the traffic to BIG-IP software. If the hardware DoS support does not exist software uses default-internal-rate-limit to limit the good traffic (most of them are flood) to external servers. Bad packets are always dropped.

If the rate limit value is infinite the rate limit is disabled. The default value is 100000.

`detection-threshold-percent`

This parameter specifies relative threshold that uses dynamically learned 1-hour average rate to detect attacks. If the current rate (1-minute average) increases the specified percent over the 1-hour average rate, attack is detected.

If the threshold value is infinite the detection is disabled. The default value is 500.

`detection-threshold-pps`

This parameter specifies absolute threshold value. If the current rate (1-minute average) is equal or above the threshold value, attack is detected.

If the threshold value is infinite the detection is disabled. The default value is 100000.

`enforce`

This option is deprecated in version 13.1.0 and is replaced by state. Enable or disable the packet drop action of DOS detection for this attack type.

`floor`

Option to set a minimum value ("floor") for the detection-threshold-pps for this vector. The range is from 0 (no-floor) to infinity (no-detection). The default value is 5000.

`multiplier-mitigation-percentage`

Specifies the mitigation multiplier value of this specific vector in percentage when using manual-multiplier-mitigation mode, The default value used is inherited from the network dos profile.

`packet-types`

This parameter is used to specify type of packets that will be classified as Sweep/Flood attacks. There are various types of packet types that can be specified. The default value is none.

`per-dst-ip-detection-pps`

Specifies the attack detection threshold (pps) per destination IP. The default value is infinite.

`per-dst-ip-limit-pps`

Specifies the attack mitigation threshold (pps) per destination IP. The default value is infinite.

`per-source-ip-detection-pps`

Specifies the attack detection threshold (pps) per source IP. The default value is infinite.

`per-source-ip-limit-pps`

Specifies the attack mitigation threshold (pps) per source IP. The default value is infinite.

`scrubbing-category`

Specifies per-DstIP scrubbing category. The default value is none.

`scrubbing-detection-seconds`

Specifies duration in seconds for which the destination IP has been offended/attacked. The default value is 10.

`scrubbing-duration`

Specifies duration in seconds for which this IP should be scrubbed. The default value is 900.

`simulate-auto-threshold`

Option to enable/disable auto-threshold simulation by generating logs if auto-threshold based detection/mitigation would have kicked in. Only valid in manual mode. The default value is disabled.

`state`

Specifies the run time state of this signature. The default value is mitigate.

The options are:

`disabled`

Do not learn, do not collect stats.

`learn-only`

Learn/Collect stats, but do not "detect" ("alarm" in ASM-speak) any attacks,

`detect-only`

Learn/Collect stats/detect, but do not mitigate (rate-limit/drop, challenge, etc.) any attacks.

`mitigate`

Learn/Collect stats/detect/mitigate (using whichever mitigations are configured).

suspicious

Specifies if the vector considers all packets or only unsolicited packets. The default value is false.

threshold-mode

Enables the threshold mode for DoS detection and DoS mitigation. The default value is manual.

The options are:

manual

Specifies the manual thresholds.

stress-based-mitigation

Specifies the manual detection ("alarm") threshold, but mitigation threshold is stress-based.

fully-automatic

Specifies both the detection ("alarm") and mitigation thresholds are automatically computed.

manual-multiplier-mitigation

Specifies the detection ("alarm") threshold is automatically computed. The mitigation threshold is calculated by the detection threshold multiplies the multiplier-mitigation-percentage.

valid-domains

Adds, deletes, modifies, or replaces a set of valid fully qualified domain names (FQDNs).

SEE ALSO

list, modify, security, security dos, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2019-07-24 security dos device-config(1)

security dos dns-nxdomain-stat

NAME

dns-nxdomain-stat - Displays and resets dos dns-nxdomain statistics of the specified context on the BIG-IP(r) system.

MODULE

security dos

SYNTAX

Manage the virtual component within the security dos module using the syntax in the following sections.

DISPLAY

show dns-nxdomain-stat

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

MODIFY

reset-stats dns-nxdomain-stat

DESCRIPTION

You can use the dns-nxdomain-stat component to display or reset dos dns-nxdomain statistics.

EXAMPLES

show dns-nxdomain-stat

Displays the dns-nxdomain's statistics in the system default units.

show dns-nxdomain-stat raw

Displays the raw dns-nxdomain's statistics.

reset-stats dns-nxdomain-stat

Resets the dns-nxdomain statistics.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2017. All rights reserved.

BIG-IP 2017-12-14 security dos dns-nxdomain-stat(1)

security dos dos-signature

NAME

dos-signature - Configures DoS Behavioral Signature(s).

MODULE

security dos

SYNTAX

Configure the dos-signature component within the security dos module using the syntax shown in the following sections.

CREATE/MODIFY

create dos-signature [name]

modify dos-signature [name]

options:

alias [string]

app-service [string | none]

approval-state [unapproved | manually-approved]

parent-context-type [device | virtual-server | device-netflow]

parent-context [string]

parent-profile [string]

description [string]

family [dns| network | http | tls]

hardware-offload [disabled | enabled]

manual-detection-threshold [integer]

manual-mitigation-threshold [integer]

multiplier-mitigation-percentage [integer]

origin [dynamic-bdos | user-defined]

predicates [list of struct(string, string, string)]

shareability-state [not-shareable | fully-shareable]

state [disabled | learn-only | detect-only | mitigate]

tags [list of string]

threshold-mode [manual | manual-multiplier-mitigation | stress-based-mitigation | fully-automatic]

type [dynamic | persistent]

DISPLAY

list dos-signature [name]

DELETE

delete dos-signature [name]

DESCRIPTION

You can use the dos-signature component to modify or display a DoS signature.

EXAMPLES

```
create security dos dos-signature Sig_Device_ToS type persistent family http origin user-defined state disabled
```

This example shows how to create a DoS signature named Sig_Device_ToS

```
list security dos dos-signature Sig_Device_ToS
```

This example shows how to display a DoS signature named Sig_Device_ToS

```
modify dos-signature Sig_Device_TTL manual-detection-threshold 10000 manual-mitigation-threshold 4294967295
```

This examples show how to modify the manual detection and mitigation threshold of a DoS signature named Sig_Device_TTL

```
delete security dos dos-signature Sig_Device_ToS
```

This example shows how to delete a DoS signature named Sig_Device_ToS

OPTIONS

alias

Specifies the alias name of a signature. The default is empty string.

app-service

Specifies the application service that the object belongs to.

approval-state

Specifies whether or not the signature has been reviewed for quality/correctness. For a persistent signature with dns or network family, the default is manually-approved. Otherwise, the default is unapproved.

User can't modify approval-state for a dynamic signature with dns or network family.

The options are:

unapproved

Specifies the signature is not approved.

manually-approved

Specifies the signature has been reviewed for quality/correctness.

parent-context-type

Specifies the type of the context for which this signature has been generated.

The available options:

device

Specifies the context type is a DoS device.

virtual-server

Specifies the type of the context is a Virtual Server.

device-netflow

Specifies the context type is Netflow device.

For a dynamic type signature, it is required field and it is not allowed to be modified once specified.

For persistent type signature, it can't be reset once it is set. The default is unspecified.

For persistent type signature with dns or network family, this field is not applicable.

parent-context

Specifies the context for which this signature has been generated. The default is empty string.

This field is based on parent-context-type. If parent-context-type is device, it must be constant "Device". If parent-context-type is device-netflow, it must be constant "NetFlow".

For a dynamic type signature, it can't be empty and it is not allowed to be modified once specified.

For persistent type signature, it can't be reset once it is set.

For persistent type signature with dns or network family, this field is not applicable.

parent-profile

Specifies the profile for which this signature has been generated. The default is empty string.

This field is based on parent-context-type. If parent-context-type is device or device-netflow, it must be constant "/Common/dos-device-config".

For a dynamic type signature, it can't be empty and it is not allowed to be modified once specified.

For a persistent type signature, it can't be reset once it is set.

This field is required for a persistent type signature with dns or network family and non-shareable shareability-state.

description

Specifies user defined description for this signature.

family

Specifies the family this signature belongs to. This is a require field for creation. The options are dns, network, http

It is not allowed to be modified once it is created.

hardware-offload

Enables or disables hardware offloading on the dynamic and persistent network family signature. The default value is enabled.

manual-detection-threshold

Specifies the manual threshold (Events Per Second) above which the traffic is declared as an attack. The default is infinite(4294967295).

This field is taken effective only when threshold-mode attribute is set to manual. For a signature with http family, it should be always 0.

For a persistent signature with dns or network family, this field is not applicable and it should be always default value.

For a dynamic signature with dns or network family, this field can't be changed if threshold-mode is fully-automatic.

manual-mitigation-threshold

Specifies the manual threshold (Events Per Second) above which the system rate limits (drops) the traffic

that matches this signature. The default is infinite(4294967295).

This field is taken effective only when threshold-mode attribute is set to manual. For a signature with http family, it should be always 0.

For a persistent signature with dns or network family, this field is not applicable and it should be always default value.

For a dynamic signature with dns or network family, this field can't be changed if threshold-mode is fully-automatic.

For a signature with parent-context-type is device-netflow, this field must be infinite(4294967295).

multiplier-mitigation-percentage
Specifies the mitigation multiplier value of this specific dos signature in percentage when using manual-multiplier-mitigation mode. The default value is inherited from the corresponding device level/profile mitigation multiplier value of the same dos family.

origin
Specifies the origin where this signature is generated from. The options are dynamic-bdos and user-defined. The default is user-defined.

It is not allowed to be modified once it is created.

predicates
Specifies list of predicates that constitutes this signature. Each predicate contains 3 string fields: metric, operator, and arguments. It is required field.

User can't add/modify predicates for a dynamic signature with dns or network family.

shareability-state
Specifies whether or not the signature can be used by Contexts (Virtual Servers) other than the one that created the signature. For a persistent signature with dns or network, the default is fully-shareable. Otherwise, the default is not-shareable.

User can't modify shareability-state for a dynamic signature with dns or network family.

This field can't be changed from fully-shareable to not-shareable if the signature is referred.

The options are:

not-shareable
Specifies the signature can only be used by context which created it.

fully-shareable
Specifies the signature can be used by contexts other than the one that created it.

state
Specifies the deployment state of this signature. The default is disabled.

The options are:

disabled
Do not learn, do not collect stats.

learn-only
Learn/Collect stats, but do not "detect" ("alarm" in ASM-speak) any attacks,

detect-only
Learn/Collect stats/detect, but do not mitigate (rate-limit/drop, challenge, etc.) any attacks.

mitigate
Learn/Collect stats/detect/mitigate (using whichever mitigation(s) are configured).

For a persistent signature with dns or network family, this field is not applicable and it should be always default value.

For a dynamic signature with dns or network family, learn-only is not allowed.

For a signature with http family, only learn-only or mitigate is allowed.

tags Specifies list of tags of this signature. The default is empty.

threshold-mode
Specifies the threshold mode for DoS detection and mitigation. The default is manual.

The options are:

manual
Specifies the manual thresholds.

stress-based-mitigation
Specifies the manual detection ("alarm") threshold, but mitigation threshold is stress-based. This option is not available for a signature with http family or for a signature with parent-context-type being device-netflow.

fully-automatic

Specifies both the detection ("alarm") and mitigation thresholds are automatically computed. This option is not available for a signature with http family.

`manual-multiplier-mitigation`
Specifies the detection ("alarm") threshold is automatically computed. The mitigation threshold is calculated by the detection threshold multiplies the multiplier-mitigation-percentage.

For a persistent signature with dns or network family, this field is not applicable and it should be always default value.

For a signature with parent-context-type is device-netflow, this field can't be stress-based-mitigation.

For a signature with http family, this field can't be stress-based-mitigation or fully-automatic.

`type` Specifies the type of this signature. The options are dynamic and persistent. The default is persistent.

It is not allowed to be changed from persistent to dynamic. User can't create dynamic signature but can modify and delete it.

SEE ALSO

`edit`, `list`, `modify`, `security`, `security dos`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2019-05-21 security dos dos-signature(1)

security dos dynamic-signatures

NAME

`dynamic-signatures` - Configures the dynamic signature(s) generated by L4 BDoS (or Dynamic Signature) AFM feature based on traffic characterization and anomaly detection.

This component has been deprecated and replaced by `dos-signature` in 13.1.0.

MODULE

`security dos`

SYNTAX

Configure the `dynamic-signatures` component within the `security dos` module using the syntax shown in the following sections.

CREATE

Currently this option is not supported for dynamic signatures.

MODIFY

`modify dynamic-signatures [name]`

options:

`context-name [name]`
`detection-threshold [integer]`
`dynamic-vectors`
`enforce [disabled | enabled]`
`mitigation-threshold [integer]`
`partition [name]`
`status [disabled | enabled]`

DISPLAY

`list dynamic-signatures`

DELETE

Currently this option is not supported for dynamic signatures.

DESCRIPTION

You can use the `dynamic-signatures` component to modify or display a dynamic signature.

EXAMPLES

`modify dynamic-signatures Sig_Device_ToS status disabled`

This example shows how to disable a dynamic signature named `Sig_Device_ToS`

`modify dynamic-signatures Sig_Device_TTL detection-threshold 10000 mitigation-threshold 4294967295`

This examples show how to modify the detection and mitigation threshold of a dynamic signature named `Sig_Device_TTL`

OPTIONS

context-name

Specifies the context for which the dynamic signature has been generated. This is a read-only field and possible values are 'Device' or 'Virtual server Name'.

detection-threshold

Specifies the threshold value above which the traffic is declared as 'anomalous' (or an attack). When the system generates a dynamic signature (based on traffic anomaly characterization), it assigns a value for detection threshold (based on various factors such as sensitivity, anomaly percent, confidence level etc.)

User can override this value by modifying the signature and specify a new value to be used for detection mechanism.

dynamic-vectors

Specifies the list of metrics and the corresponding values/ranges that constitutes a dynamic signature. This is a read-only field.

enforce

Specifies the run time behavior of the dynamic signature in the datapath with respect to enforcement. Possible values are: disabled or enabled.

If set to disabled, the system does not enforce the signature to rate limit traffic but only collect statistics. If set to enabled, in addition to collecting stats, system also enforces the signature to detect an attack and limit traffic as per the mitigation threshold.

mitigation-threshold

Specifies the threshold above which the system rate limits (drops) the traffic that matches this generated dynamic signature. When the system generates a dynamic signature, it assigns a value for mitigation threshold based on certain factors such as mitigation configuration, detection threshold etc.

User can override this value by modifying the signature and specify a new value to be used for mitigating traffic that matches this dynamic signature.

status

Specifies the run time status of the generated signature. Possible values are: disabled or enabled.

By default, the status is set to enabled when the system generates a dynamic signature. User can disable detection and mitigation for this dynamic signature by setting this field to disabled.

SEE ALSO

edit, list, modify, security, security dos, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-03-09 security dos dynamic-signatures(1)

security dos ip-uncommon-protolist

NAME

ip-uncommon-protolist - Configures the DoS IP uncommon protocol list component within the security dos module using the syntax shown in the following sections.

MODULE

security dos

SYNTAX

CREATE/MODIFY

modify ip-uncommon-protolist [name]

options:

description [string]

entries [add | delete | none | replace-all-with]

DISPLAY

list ip-uncommon-protolist

list ip-uncommon-protolist [name]

DESCRIPTION

You can use the ip-uncommon-protolist component to configure a DoS IP uncommon protocol list. These DoS IP uncommon protocol entries are used to rate limit the PPS (Packet per second) rate for IP protocols (identified by the IP protocol number) that are not expected to be common in the customer deployment.

EXAMPLES

create ip-uncommon-protolist plistA entries add { 20 tcp } Creates a plistA IP uncommon protocol list with

entries protocol 20 and tcp.

modify ip-uncommon-protolist plistA entries add { udp } Modifies the plistA to add udp protocol to the list.

modify ip-uncommon-protolist plistA entries delete { 20 } Modifies the plistA to add protocol 20 from the list.

list ip-uncommon-protolist plistA Displays the current list of plistA entries.

OPTIONS

description

Specified user defined description for a DoS ip-uncommon--protolist.

entries

Specified contexts of a DoS ip-uncommon--protolist. The default is none.

The options are:

add Add a protocol entry to the specified IP uncommon protocol list.

delete

Remove a protocol entry from the specified IP uncommon protocol list.

none Remove all protocols from the specified IP uncommon protocol list.

replace-all-with

Replace all protocols in the specified IP uncommon protocol list.

name Specified name of a DoS ip-uncommon--protolist. This is required during creation.

"ip-uncommon-protolist" is the system ip-uncommon--protolist. It can't be deleted.

SEE ALSO

create, edit, list, modify, security, security dos, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017. All rights reserved.

BIG-IP 2017-05-01 security dos ip-uncommon-protolist(1)

security dos l4bdos-file-object

NAME

l4bdos-file-object - L4 BDoS internal file object.

MODULE

security

SYNTAX

Configure l4bdos-file-object within security dos module using the syntax shown in the following sections.

CREATE/MODIFY

create l4bdos-file-object [name]

modify l4bdos-file-object [name]

options:

app-service [[string] | none]

context-name [string]

source-path [file name]

DELETE

delete l4bdos-file-object [name]

DESCRIPTION

l4bdos file object is used internally for synchronization of bdos module data across HA (high availability) or Cluster setup.

EXAMPLES

create name source-path /var/bdos/file.json context-name /Common/VS1

Creates a file object named name, that gets its contents from file object /var/bdos/file.json and is created for ltm virtual /Common/VS1.

delete filename

Deletes the file object named filename.

OPTIONS

app-service

Specifies the name of the application service to which the file object belongs. The default value is none.

source-path

The name of the source file from which the contents are copied to create the file object.

context-name

The name of the virtual server to which this file object belongs.

SEE ALSO

create, delete, edit, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2017-07-07 security dos l4bdos-file-object(1)

security dos network-whitelist

NAME

network-whitelist - Configures the DoS network whitelist component within the security dos module using the syntax shown in the following sections. These DoS network whitelist entries are applied to all packets except those going through the management interface.

MODULE

security dos

SYNTAX

MODIFY

```
modify network-whitelist dos-network-whitelist
options:
  address-list [name]
  description [string]
  entries [add | delete | modify | replace-all-with] {
    [ [name] ] {
      options:
    }
  }
  description [string]
  destination {
    address [ip_address/prefixlen]
    port [port]
  }
  ip-protocol [any | icmp | igmp | tcp | udp]
  match-ip-version [false | true]
  source {
    address [ip_address/prefixlen] ]
    vlans [vlan name | vlanid/mask]
  }
}
entries none
extended-entries [add | delete | modify | replace-all-with] {
  [ [name] ] {
    options:
  }
  description [string]
  destination {
    address [ip_address/prefixlen]
    port [port]
  }
  ip-protocol [any | icmp | igmp | tcp | udp]
  match-ip-version [false | true]
  source {
    address [ip_address/prefixlen] ]
    vlans [vlan name | vlanid/mask]
  }
}
extended-entries none
```

DISPLAY

```
list network-whitelist
```

DESCRIPTION

You can use the network-whitelist component to configure two types of DoS network whitelists: 1) standard whitelist, up to eight entries; 2) extended whitelist, up to the the number of entries specified by DB variable dos.maxewlsize (range from 0 to 1024). Whitelists configured this way can be applied to all traffic except those from the management interface. Along with that you can use address-list to configure the srcIP Global whitelist. To this address-list you need to attach the address list objects. This address-list can be a nested list of fully qualified address. Subnets and IP address ranges and geo-locations are not allowed. The HSB hardware compares all incoming traffic to the network-whitelist entries. If a match is found then it does not do DoS vector checks for those packets. If a match is not found then DoS vector checks are done on those packets. The network software does its regular DoS vector checks on the incoming packets as usual. If a DoS vector is hit then it compares that packet with the DoS network-whitelist entries. If the packet matches an entry, then the system does not increment the DoS vector that matched. If the packets does not match a DoS network-whitelist entry then the matched DoS vector is incremented and appropriate action is taken.

If an entry specifies more than one of the above items, a packet must pass all of the items to successfully match. For example, if an entry specifies a source subnet and a destination port, a packet must originate from the given subnet and must also have the specified destination port.

Either destination ip_address/prefixlen or source ip_address/prefixlen can be specified in a network-whitelist entry. An ip_address/prefixlen for both source and destination cannot be specified for an entry.

EXAMPLES

```
modify network-whitelist dos-network-whitelist description "bad interfaces" entries add { re_telnet { ip-protocol tcp destination { port telnet } } }
```

Creates a new entry called re_telnet. It matches any TCP packet whose destination port is telnet.

```
modify network-whitelist dos-network-whitelist entries add { internal-net { source { address 172.27.0.0/16 } } }
```

Creates an entry that matches traffic from the 172.27.0.0 network.

```
list network-whitelist
security dos network-whitelist dos-network-whitelist {
  entries {
    re_telnet {
      ip-protocol tcp
      destination {
        port telnet
      }
    }
    internal-net {
      source {
        address 172.27.0.0/16
      }
    }
  }
}
```

Displays the current list of DoS whitelist entries.

```
modify network-whitelist dos-network-whitelist entries delete { internal-net }
```

Removes the "internal-net" entry from the list of network-whitelist entries.

```
modify security dos network-whitelist dos-network-whitelist extended-entries add { netwl { source { address 10.0.0.0/8 } destination { address 20.20.20.0/24 } ip-protocol udp } }
```

Creates a new extended entry called netwl. It matches any UDP packet matches source network address 10.x.x.x and destination network address 20.20.20.x.

```
list security dos network-whitelist dos-network-whitelist extended-entries { netwl }
security dos network-whitelist dos-network-whitelist {
  extended-entries {
    netwl {
      description none
      ip-protocol udp
      destination {
        address 20.20.20.0/24
      }
      port any
    }
    source {
      address 10.0.0.0/8
    }
    vlans any
  }
}
```

Displays the extended whitelist entry just configured.

OPTIONS

address-list
Specifies the object in security firewall address-list as the srcIP Global whitelist.

description
Your description for the DoS network-whitelist entries.

entries

Adds, deletes, or replaces a standard network-whitelist entry, by specifying an entry name. If an entry by the specified name does not exist, it will be created.

add Creates a new entry, which you specify next with a unique string in curly braces ({}).

delete

Deletes the entry that you specify next, in curly braces ({}). You can use `delete {all}` to empty the list of network-whitelist entries, which has the same effect as using `none` (see below).

modify

Modifies the existing entry that you specify next, in curly braces ({}). After the entry name, enter the new configuration settings for the entry inside a nested set of curly braces.

replace-all-with

Replaces the current set of network-whitelist entries with the entry(s) that you specify next, in curly braces ({}).

none Empties the list of network-whitelist entries.

Enter the name of a entry to be added or modified, then enter an open curly brace ({}), one or more of the following options, and a closed curly brace ({}).

description

Your description for the current entry.

destination

Matches against each packet's destination IP and/or destination port.

address

Specifies an IP address and network to compare against the packet's destination address.

The format for an IPv4 address is `a.b.c.d[/prefix]`. The general format for an IPv6 address is `a:b:c:d:e:f:g:h[/prefix]`; you can shorten this by eliminating leading zeros from each field (for example, you can shorten "2001:0db7:3f4a:09dd:ca90:ff00:0042:8329" to "2001:db7:3f4a:9dd:ca90:ff00:42:8329"), and/or by removing the longest contiguous field of zeros (for example, you can shorten "2001:0:0:0:c34a:0:23ff:678" to "2001::c34a:0:23ff:678"). TMSH accepts any valid text representation of IPv6 addresses, as defined in RFC 2373 (see).

port Specifies a port to compare against the packet's destination port.

ip-protocol

Specifies the IP protocol to compare against the packet. This could be any, icmp, igmp, tcp or udp. If you specify this option, a packet only matches if it uses the chosen protocol.

match-ip-version

Specifies whether any/any6 in source/destination address is to match 'any IPv4', 'any IPv6', or 'any IPv4 and any IPv6' addresses. If `match-ip-version` is true, both source and destination addresses must have the same IP address family. If `match-ip-version` is false and both source and destination addresses are any or any6, both addresses represent 'any IPv4 and IPv6 addresses'. If `match-ip-version` is false and only one address is set to any or any6, the address is interpreted based on the other IP address' family (IPv4 or IPv6). The default is false.

source

Matches against each packet's source IP, and/or source VLANs.

address

Specifies an IP address and network to compare against the packet's source address.

The format for an IPv4 address is `a.b.c.d`. The general format for an IPv6 address is `a:b:c:d:e:f:g:h`.

vlan

Specifies either a vlan name or a range of vlanids to compare against the packet. The range is specified as `vlanid/mask`. For example if you specify "3200/8" then the vlanid range will be 3200-3327.

extended-entries

Adds, deletes, or replaces an extended network-whitelist entry, by specifying an entry name. If an entry by the specified name does not exist, it will be created.

add Creates a new entry, which you specify next with a unique string in curly braces ({}).

delete

Deletes the entry that you specify next, in curly braces ({}). You can use `delete {all}` to empty the list of network-whitelist entries, which has the same effect as using `none` (see below).

modify

Modifies the existing entry that you specify next, in curly braces ({}). After the entry name, enter the new configuration settings for the entry inside a nested set of curly braces.

replace-all-with

Replaces the current set of network-whitelist entries with the entry(s) that you specify next, in curly braces ({}).

none Empties the list of network-whitelist extended-entries.

Enter the name of a entry to be added or modified, then enter an open curly brace ({}), one or more of the following options, and a closed curly brace (}).

description

Your description for the current entry.

destination

Matches against each packet's destination IP and/or destination port.

address

Specifies an IP address and network to compare against the packet's destination address.

The format for an IPv4 address is a.b.c.d[/prefix]. The general format for an IPv6 address is a:b:c:d:e:f:g:h[/prefix]; you can shorten this by eliminating leading zeros from each field (for example, you can shorten "2001:0db7:3f4a:09dd:ca90:ff00:0042:8329" to "2001:db7:3f4a:9dd:ca90:ff00:42:8329"), and/or by removing the longest contiguous field of zeros (for example, you can shorten "2001:0:0:0:c34a:0:23ff:678" to "2001::c34a:0:23ff:678"). TMSH accepts any valid text representation of IPv6 addresses, as defined in RFC 2373 (see).

port Specifies a port to compare against the packet's destination port.

ip-protocol

Specifies the IP protocol to compare against the packet. This could be any, icmp, igmp, tcp or udp. If you specify this option, a packet only matches if it uses the chosen protocol.

match-ip-version

Specifies whether any/any6 in source/destination address is to match 'any IPv4', 'any IPv6', or 'any IPv4 and any IPv6' addresses. If match-ip-version is true, both source and destination addresses must have the same IP address family. If match-ip-version is false and both source and destination addresses are any or any6, both addresses represent 'any IPv4 and IPv6 addresses'. If match-ip-version is false and only one address is set to any or any6, the address is interpreted based on the other IP address' family (IPv4 or IPv6). The default is false.

source

Matches against each packet's source IP, and/or source VLANs.

address

Specifies an IP address and network to compare against the packet's source address.

The format for an IPv4 address is a.b.c.d. The general format for an IPv6 address is a:b:c:d:e:f:g:h.

vlan

Specifies either a vlan name or a range of vlanids to compare against the packet. The range is specified as vlanid/mask. For example if you specify "3200/8" then the vlanid range will be 3200-3327.

EXAMPLES

```
modify security dos network-whitelist dos-network-whitelist address-list [name]
```

It adds list1 objects to the global address-list. For configuring the address list objects (list1) you can use the following examples:

```
create security firewall address-list list1 addresses [add | delete] { 30.30.30.30 45:56:567:234:456::0 }
```

```
list security firewall address-list list1
```

```
security firewall address-list list1 {  
  addresses {  
    30.30.30.30 { }  
    45:56:567:234:456:: { }  
  }  
}
```

This is how you can list the address-list objects that you configured for global whitelists list security dos network-whitelist address-list security dos network-whitelist dos-network-whitelist { address-list list1 }

SEE ALSO

edit, list, modify, security, security dos, tms security firewall address-lists

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2016. All rights reserved.

BIG-IP 2018-03-26 security dos network-whitelist(1)

security dos profile

NAME

profile - Configures a DoS profile.

MODULE

security dos

SYNTAX

Configure the profile component within the security dos module using the syntax shown in the following sections.

CREATE/MODIFY

create profile [name]

modify profile [name]

options:

app-service [[string] | none]

application [none | add | delete | modify | replace-all-with] {

[sub-profile-name] ... {

options:

bot-defense {

collect-stats [enabled | disabled]

cross-domain-requests [allow-all | validate-bulk | validate-upon-request]

external-domains [none | add | delete | replace-all-with] { [string] ... }

grace-period [integer]

mode [always | disabled | during-attacks]

site-domains [none | add | delete | replace-all-with] { [string] ... }

url-whitelist [none | add | delete | replace-all-with] { [string] ... }

browser-legit-enabled [enabled | disabled]

browser-legit-captcha [enabled | disabled]

}

bot-signatures {

categories [none | add | delete | modify | replace-all-with] {

action {

[block | none | report]

}

}

check [enabled | disabled]

disabled-signatures [none | add | delete | modify | replace-all-with]

}

captcha-response {

failure {

body [string]

type [custom | default]

}

first {

body [string]

type [custom | default]

}

}

geolocations [none | add | delete | modify | replace-all-with] {

options:

[black-listed | white-listed]

}

heavy-urls {

automatic-detection [enabled | disabled]

exclude [none | add | delete | replace-all-with] { [string] ... }

include [none | add | delete | replace-all-with] { [string] ... }

include-list [none | add | delete | replace-all-with] { [string] { [integer] } ... }

latency-threshold [integer]

protection [enabled | disabled]

}

ip-whitelist [none | add | delete | modify | replace-all-with] {

[address ... | address/mask ...]

}

stress-based {

de-escalation-period [integer]

escalation-period [integer]

geo-captcha-challenge [enabled | disabled]

geo-client-side-defense [enabled | disabled]

geo-minimum-share [integer]

geo-rate-limiting [enabled | disabled]

geo-request-blocking-mode [block-all | rate-limit]

geo-share-increase-rate [integer]

geo-maximum-auto-tps [integer]

geo-minimum-auto-tps [integer]

ip-captcha-challenge [enabled | disabled]

ip-client-side-defense [enabled | disabled]

ip-maximum-tps [integer]

ip-minimum-tps [integer]

ip-rate-limiting [enabled | disabled]

ip-request-blocking-mode [block-all | rate-limit]

ip-tps-increase-rate [integer]

ip-maximum-auto-tps [integer]

ip-minimum-auto-tps [integer]

mode [off | transparent | blocking]

```
thresholds-mode [manual | automatic]
site-captcha-challenge [enabled | disabled]
site-client-side-defense [enabled | disabled]
site-maximum-tps [integer]
site-minimum-tps [integer]
site-rate-limiting [enabled | disabled]
site-tps-increase-rate [integer]
site-maximum-auto-tps [integer]
site-minimum-auto-tps [integer]
static-url-mitigation [enabled | disabled]
url-captcha-challenge [enabled | disabled]
url-client-side-defense [enabled | disabled]
url-maximum-tps [integer]
url-minimum-tps [integer]
url-rate-limiting [enabled | disabled]
url-tps-increase-rate [integer]
url-maximum-auto-tps [integer]
url-minimum-auto-tps [integer]
url-enable-heavy [enabled | disabled]
device-captcha-challenge [enabled | disabled]
device-client-side-defense [enabled | disabled]
device-maximum-tps [integer]
device-minimum-tps [integer]
device-rate-limiting [enabled | disabled]
device-request-blocking-mode [block-all | rate-limit]
device-tps-increase-rate [integer]
device-maximum-auto-tps [integer]
device-minimum-auto-tps [integer]
behavioral {
  dos-detection [enabled | disabled]
  mitigation-mode [none | conservative | standard | aggressive ]
  signatures [enabled | disabled]
  signatures-approved-only [disabled | disabled]
  accelerated-signatures [enables | disabled]
  tls-signatures [enabled | disabled]
  tls-fp [enabled | disabled]
}
}
tcp-dump {
  maximum-duration [integer]
  maximum-size [integer]
  record-traffic [enabled | disabled]
  repetition-interval [[integer] | once-per-attack]
}
tps-based {
  de-escalation-period [integer]
  escalation-period [integer]
  geo-captcha-challenge [enabled | disabled]
  geo-client-side-defense [enabled | disabled]
  geo-minimum-share [integer]
  geo-rate-limiting [enabled | disabled]
  geo-request-blocking-mode [block-all | rate-limit]
  geo-share-increase-rate [integer]
  ip-captcha-challenge [enabled | disabled]
  ip-client-side-defense [enabled | disabled]
  ip-maximum-tps [integer]
  ip-minimum-tps [integer]
  ip-rate-limiting [enabled | disabled]
  ip-request-blocking-mode [block-all | rate-limit]
  ip-tps-increase-rate [integer]
  ip-maximum-auto-tps [integer]
  ip-minimum-auto-tps [integer]
  mode [off | transparent | blocking]
  thresholds-mode [manual | automatic]
  site-captcha-challenge [enabled | disabled]
  site-client-side-defense [enabled | disabled]
  site-maximum-tps [integer]
  site-minimum-tps [integer]
  site-rate-limiting [enabled | disabled]
  site-tps-increase-rate [integer]
  site-maximum-auto-tps [integer]
  site-minimum-auto-tps [integer]
  static-url-mitigation [enabled | disabled]
  url-captcha-challenge [enabled | disabled]
  url-client-side-defense [enabled | disabled]
  url-maximum-tps [integer]
  url-minimum-tps [integer]
  url-rate-limiting [enabled | disabled]
  url-tps-increase-rate [integer]
  url-maximum-auto-tps [integer]
  url-minimum-auto-tps [integer]
  url-enable-heavy [enabled | disabled]
  device-captcha-challenge [enabled | disabled]
  device-client-side-defense [enabled | disabled]
  device-maximum-tps [integer]
  device-minimum-tps [integer]
  device-rate-limiting [enabled | disabled]
```

```

device-request-blocking-mode [block-all | rate-limit]
device-tps-increase-rate [integer]
device-maximum-auto-tps [integer]
device-minimum-auto-tps [integer]
}
trigger-irule [enabled | disabled]
single-page-application [enabled | disabled]
scrubbing-enable [enabled | disabled]
scrubbing-duration-sec [integer]
rtbh-enable [enabled | disabled]
rtbh-duration-sec [integer]
fastl4-acceleration-profile [fastL4 profile name]
}
}
custom-signatures [none | add | delete | modify | replace-all-with] {
  name [string] {
options:
manual-detection-threshold [integer]
manual-mitigation-threshold [integer]
state [detect-only | disabled | learn-only | mitigate]
threshold-mode [fully-automatic | manual | stress-based-mitigation]
}
}
description [string]
dos-network [none | add | delete | modify | replace-all-with] {
  [sub-profile-name] ... {
options:
dynamic-signatures {
  detection [disabled | enabled | learn-only]
  mitigation [none | low | medium | high | manual-multiplier]
  scrubber-advertisement-period [integer]
  scrubber-category [name]
  scrubber-enable [yes | no]
}
multiplier-mitigation-percentage [integer]
network-attack-vector [none | add | delete | modify | replace-all-with] {
  attack-type [ext-hdr-too-large | hop-cnt-low | host-unreachable |
icmpv4-flood | icmpv6-flood | icmp-frag | ip-frag-flood |
ip-opt-frames | ipv6-ext-hdr-frames | ipv6-frag-flood |
non-tcp-connection | opt-present-with-illegal-len | sweep |
tcp-half-open | tcp-opt-overruns-tcp-hdr | tcp-psh-flood |
tcp-rst-flood | tcp-syn-flood | tcp-synack-flood | tcp-syn-oversize |
tcp-bad-urg | tcp-window-size | tidcmp | too-many-ext-hdrs |
udp-flood | unk-tcp-opt-type]
options:
enforce [disabled | enabled]
auto-blacklisting [disabled | enabled]
auto-threshold [disabled | enabled ]
allow-upstream-scrubbing [disabled | enabled]
attacked-dst [disabled | enabled]
auto-scrubbing [disabled | enabled]
bad-actor [disabled | enabled]
blacklist-detection-seconds [integer]
blacklist-duration [integer]
blacklist-category [enter name of ip-intelligence category]
multiplier-mitigation-percentage [integer]
per-source-ip-detection-pps [integer]
per-source-ip-limit-pps [integer]
per-dst-ip-detection-pps [integer]
per-dst-ip-limit-pps [integer]
scrubbing-category [[category name] | none]
scrubbing-detection-seconds [integer]
scrubbing-duration [integer]
rate-increase [integer]
rate-limit [integer | infinite]
rate-threshold [integer | infinite]
packet-types [suspicious | ipfrag | exthdr | tcp-syn-only |
tcp-synack | tcp-rst | host-unrch | tidcmp | icmp | udp-flood |
dns-query-a | dns-query-aaaa | dns-query-any | dns-query-axfr |
dns-query-cname | dns-query-ixfr | dns-query-mx | dns-query-ns
| dns-query-other | dns-query-ptr | dns-query-soa |
dns-query-srv | dns-query-src | dns-query-txt | sip-method-ack
| sip-method-cancel | sip-method-message | sip-method-options |
sip-method-prack | sip-method-register | sip-method-bye |
sip-method-invite | sip-method-notify | sip-method-other |
sip-method-publish | sip-method-subscribe ]
state [disabled | learn-only | detect-only | mitigate]
suspicious [ false | true ]
threshold-mode [manual | stress-based-mitigation | fully-automatic]
}
}
}
protocol-dns [none | add | delete | modify | replace-all-with] {
  [sub-profile-name] ... {
options:
dns-query-vector [none | add | delete | modify | replace-all-with] {
  query-type [a | aaaa | any | axfr | cname | ixfr | mx | ns | nxdomain |

```

```

other | ptr | soa | srv | txt ]
options:
  enforce [disabled | enabled]
  auto-blacklisting [disabled | enabled]
  auto-threshold [disabled | enabled ]
  allow-upstream-scrubbing [disabled | enabled]
  attacked-dst [disabled | enabled]
  auto-scrubbing [disabled | enabled]
  bad-actor [disabled | enabled]
  blacklist-detection-seconds [integer]
  blacklist-duration [integer]
  blacklist-category [enter name of ip-intelligence category]
  multiplier-mitigation-percentage [integer]
  per-source-ip-detection-pps [integer]
  per-source-ip-limit-pps [integer]
  per-dst-ip-detection-pps [integer]
  per-dst-ip-limit-pps [integer]
  scrubbing-category [[category name] | none]
  scrubbing-detection-seconds [integer]
  scrubbing-duration [integer]
  rate-increase [integer]
  rate-limit [integer | infinite]
  rate-threshold [integer | infinite]
  state [disabled | learn-only | detect-only | mitigate]
  suspicious [ false | true ]
  threshold-mode [manual | stress-based-mitigation | fully-automatic]
  valid-domains [none | add | delete ] replace-all-with {
[domain-name] ...
}
}
multiplier-mitigation-percentage [integer]
prot-err-attack-detection [integer]
prot-err-atck-rate-incr [integer]
}
}
protocol-sip [none | add | delete | modify | replace-all-with] {
[sub-profile-name] ... {
options:
multiplier-mitigation-percentage [integer]
prot-err-atck-rate-increase [integer]
prot-err-atck-rate-threshold [integer]
prot-err-attack-detection [integer]
sip-attack-vector [none | add | delete | modify | replace-all-with] {
  type [ack | cancel | message | options | prack | register
| bye | invite | notify | other | publish | subscribe | uri-limit]
options:
  enforce [disabled | enabled]
  auto-blacklisting [disabled | enabled]
  auto-threshold [disabled | enabled ]
  allow-upstream-scrubbing [disabled | enabled]
  attacked-dst [disabled | enabled]
  auto-scrubbing [disabled | enabled]
  bad-actor [disabled | enabled]
  blacklist-detection-seconds [integer]
  blacklist-duration [integer]
  blacklist-category [enter name of ip-intelligence category]
  multiplier-mitigation-percentage [integer]
  per-source-ip-detection-pps [integer]
  per-source-ip-limit-pps [integer]
  per-dst-ip-detection-pps [integer]
  per-dst-ip-limit-pps [integer]
  scrubbing-category [[category name] | none]
  scrubbing-detection-seconds [integer]
  scrubbing-duration [integer]
  rate-increase [integer]
  rate-limit [integer | infinite]
  rate-threshold [integer | infinite]
  state [disabled | learn-only | detect-only | mitigate]
  suspicious [ false | true ]
  threshold-mode [manual | manual-multiplier-mitigation | stress-based-mitigation | fully-automatic]
}
}
}
whitelist [enter addresses list name]
http-whitelist [enter addresses list name]

reset-stats profile [ [ [name] | [glob] | [regex] ] ... ]
options:
  dos-dnsnxdomain-stat

edit profile [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties

DISPLAY
list profile

```

```
list profile [ [name] | [glob] | [regex] ] ... ]
show running-config profile
show running-config profile [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  recursive
show profile [ [name] | [glob] | [regex] ] ... ]
options:
  dns-nxdomain-stat
  field-fmt
```

```
DELETE
delete profile [name]
```

DESCRIPTION

You can use the profile component to create, modify, display, or delete a DoS profile for use with DoS Protection functionality.

EXAMPLES

```
create profile my_dos_profile
```

Creates a custom DoS profile named my_dos_profile with initial settings.

```
list profile
```

Displays the properties of all DoS profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

application

Adds, deletes, or replaces a single Application Security sub-profile. You can configure the following options for Application Security:

bot-defense

Specifies properties of proactive bot defense in Application Security. You can configure the following options for Proactive Bot Defense:

collect-stats

Enables or disables domain statistics collection.

cross-domain-requests

Specifies a cross-domain requests handling mode. The options are:

allow-all

Allows all cross-domain requests. This is the default value.

validate-bulk

System validates domains in bulk: the cookies for the related domains are created together with the cookie for the current domain, by generating challenges in iframes - one per each domain.

validate-upon-request

System validates domains upon request: the cookie for the related domain is generated when a request arrives to an unqualified URL without a cookie.

external-domains

Configures a list of external domains that are allowed to link to resources of this website.

grace-period

Specifies the length of grace period (in seconds) in which only the Simple Bot Prevention is enforced.

mode Specifies a mode of proactive bot defense. The options are:

always

Specifies that the proactive bot defense is always enabled.

disabled

Specifies that the proactive bot defense is disabled. This is the default value.

during-attacks

Specifies that the proactive bot defense is enabled only during attacks.

site-domains

Configures a list of domains that are part of the website.

url-whitelist

Configures a list of URLs to exclude from the proactive bot defense.

browser-legit-enabled

Enables or disables the proactive bot defense validation of browser legitimacy and blocking of requests from suspicious clients.

browser-legit-captcha

Enables or disables the browser legitimacy detection improvement using CAPTCHA. In order to enable it, you must first enable browser-legit-enabled.

bot-signatures

Specifies settings of Bot Signatures in Application Security. You can configure the following options for Bot Signatures:

categories

Specifies the action for each Bot Signature Category. You can configure the following options for each Bot Signature Category:

action

Specifies the action for the Bot Signature Category. The possible actions are none, block and report.

check

Enables or disables the checking of Bot Signature, allowing bots to be detected.

disabled-categories

Configures a list of disabled Bot Signatures.

captcha-response

Specifies properties of the CAPTCHA response in Application Security. You can configure the following options for CAPTCHA Response Settings:

failure

Specifies properties of a failed CAPTCHA response. You can configure the following options for a failed CAPTCHA response:

body Configures a failed CAPTCHA response body.

type Configures a type of a failed CAPTCHA response body. You can configure the following options for a failed CAPTCHA response type:

custom

Configures a custom failed CAPTCHA response type.

default

Configures a default failed CAPTCHA response type.

first

Specifies properties of the first CAPTCHA response. You can configure the following options for the first CAPTCHA response:

body Configures the first CAPTCHA response body.

type Configures a type of the first CAPTCHA response body. You can configure the following options for the first CAPTCHA response type:

custom

Configures a custom first CAPTCHA response type.

default

Configures a default first CAPTCHA response type.

geolocations

Configures a list of blacklisted/whitelisted Geolocations. You can configure the following options for each Geolocation:

[black-listed | white-listed]

Specifies a type of Geolocation.

heavy-urls

Specifies heavy URL protection in Application Security. You can configure the following options for heavy URL protection:

automatic-detection

Enables or disables automatic heavy URL detection. In order to enable it, you must first enable protection.

exclude

Configures a list of URLs (or wildcards) to exclude from the heavy URLs.

include

(Deprecated, use include-list) Configures a list of URLs to include in the heavy URLs.

include-list

Configures a list of URLs to include in the heavy URLs.

latency-threshold

Specifies the latency threshold for automatic heavy URL detection (in milliseconds).

protection

(Deprecated, use `stress/tps.url-enable-heavy`) Enables or disables heavy URL protection. To enable it, you must additionally enable one of the following DoS URL-based prevention policy methods: `url-client-side-defense` or `url-rate-limiting`. This can be done for either `tps-based` or `stress-based` anomaly protection.

`ip-whitelist`

Attribute `ip-whitelist` is deprecated in version 13.0.0; consider using `http-whitelist` instead. Adds, deletes, or replaces a set of IP addresses and subnets in the whitelist of Application Security.

`name` Specifies a dummy name for enabled Application Security. This option is required for the operations `create`, `delete`, `modify`, and `replace-all-with`.

`stress-based`

Specifies Stress-based anomaly in Application Security. You can configure the following options for Stress-based anomaly:

`de-escalation-period`

Specifies the de-escalation period (in seconds) in Stress-based anomaly.

`escalation-period`

Specifies the escalation period (in seconds) in Stress-based anomaly.

`geo-captcha-challenge`

Enables or disables Geolocation-based CAPTCHA challenge in Stress-based anomaly.

`geo-client-side-defense`

Enables or disables Geolocation-based client side integrity defense in Stress-based anomaly.

`geo-minimum-share`

Specifies the minimum traffic share for detection in Geolocation detection criteria of Stress-based anomaly.

`geo-rate-limiting`

Enables or disables Geolocation-based rate limiting in Stress-based anomaly.

`geo-request-blocking-mode`

Specifies a Geolocation-based request blocking mode of Stress-based anomaly. The options are:

`block-all`

Specifies that the system blocks all requests from the respective Geolocation.

`rate-limit`

Specifies that the system blocks requests from the respective Geolocation based on the traffic share ratio. This is the default value.

`geo-share-increase-rate`

Specifies the percentage by which TPS increased in Geolocation detection criteria of Stress-based anomaly.

`ip-captcha-challenge`

Enables or disables Source IP-based CAPTCHA challenge in Stress-based anomaly.

`ip-client-side-defense`

Enables or disables Source IP-based client side integrity defense in Stress-based anomaly.

`ip-maximum-tps`

Specifies the amount which TPS reached in IP detection criteria of Stress-based anomaly.

`ip-minimum-tps`

Specifies the minimum TPS threshold for detection in IP detection criteria of Stress-based anomaly.

`ip-rate-limiting`

Enables or disables Source IP-based rate limiting in Stress-based anomaly.

`ip-request-blocking-mode`

Specifies a Source IP-based request blocking mode of Stress-based anomaly. The options are:

`block-all`

Specifies that the system blocks all requests from the respective Source IP address.

`rate-limit`

Specifies that the system blocks requests from the respective Source IP address based on the traffic share ratio. This is the default value.

`ip-tps-increase-rate`

Specifies the percentage by which TPS increased in IP detection criteria of Stress-based anomaly.

`mode` Specifies an operation mode of Stress-based anomaly. The options are:

`off` Specifies that the system does not check for DoS attacks. This is the default value.

`transparent`

Specifies that when the system detects an attack, it displays the attack data on the Reporting DoS Attacks screen. In transparent mode the system does not drop requests either

from the attacking IP address, or to attacked URLs.

blocking

Specifies that when the system detects an attack, in addition to displaying the attack data on the Reporting DoS Attacks screen, the system also drops either connections from the attacking IP address, or requests to attacked URLs.

site-captcha-challenge

Enables or disables Site-wide CAPTCHA challenge in Stress-based anomaly.

site-client-side-defense

Enables or disables Site-wide client side integrity defense in Stress-based anomaly.

site-maximum-tps

Specifies the amount which TPS reached in Site-wide detection criteria of Stress-based anomaly.

site-minimum-tps

Specifies the minimum TPS threshold for detection in Site-wide detection criteria of Stress-based anomaly.

site-rate-limiting

Enables or disables Site-wide rate limiting in Stress-based anomaly.

site-tps-increase-rate

Specifies the percentage by which TPS increased in Site-wide detection criteria of Stress-based anomaly.

static-url-mitigation

Enables or disables Static URL mitigation in Stress-based anomaly.

url-captcha-challenge

Enables or disables URL-based CAPTCHA challenge in Stress-based anomaly.

url-client-side-defense

Enables or disables URL-based client side integrity defense in Stress-based anomaly.

url-maximum-tps

Specifies the amount which TPS reached in URL detection criteria of Stress-based anomaly.

url-minimum-tps

Specifies the minimum TPS threshold for detection in URL detection criteria of Stress-based anomaly.

url-rate-limiting

Enables or disables URL-based rate limiting in Stress-based anomaly.

url-tps-increase-rate

Specifies the percentage by which TPS increased in URL detection criteria of Stress-based anomaly.

behavioral

Specifies properties of Behavioral Detection in Stress-based anomaly. You can configure the following options for Behavioral Detection:

dos-detection

Enables or disables the Behavior Based Detection.

mitigation-mode

Specifies mitigation impact on suspicious bad actors/requests. None: Learns and monitors traffic behavior, but no action is taken. Conservative protection:If enabled, slows down and rate limits requests from anomalous IP addresses based on its anomaly detection confidence and the server's health. If enabled, blocks requests that match the attack signatures. Standard protection:If enabled, slows down requests from anomalous IP addresses based on its anomaly detection confidence and the server's health. Rate limits requests from anomalous IP addresses and, if necessary, rate limits all requests based on the servers health. Limits the number of concurrent connections from anomalous IP addresses and, if necessary, limits the number of all concurrent connections based on the server's health. If enabled, blocks requests that match the attack signatures. Aggressive protection:If enabled, slows down requests from anomalous IP addresses based on its anomaly detection confidence and the server's health. Rate limits requests from anomalous IP addresses and, if necessary, rate limits all requests based on the servers health. Limits the number of concurrent connections from anomalous IP addresses and, if necessary, limits the number of all concurrent connections based on the server's health. Proactively performs all protection actions (even before an attack). Increases the impact of the protection techniques. If enabled, blocks requests that match the attack signatures. Increases the impact of blocked requests.

signatures

Enables or disables signature usage and mitigation.

signatures-approved-only

Allows to use only manually approved signatures.

accelerated-signatures

Enables or disables signatures detection before the connection establishment. Automatically enables syn-cookie mechanism during attack.

tls-signatures

Enables or disables TLS signatures detection before the connection establishment.

tls-fp

Enables or disables TLS patterns as an extension of bad actors detection.

tcp-dump

Specifies properties of traffic recording during attacks in Application Security. You can configure the following options for Record Traffic During Attacks:

maximum-duration

Specifies the TCP dump maximum duration (in seconds).

maximum-size

Specifies the TCP dump maximum size (in megabytes).

record-traffic

Enables or disables traffic recording during attacks.

repetition-interval

Specifies the TCP dump repetition interval (in seconds).

tps-based

Specifies TPS-based anomaly in Application Security. You can configure the following options for TPS-based anomaly:

de-escalation-period

Specifies the de-escalation period (in seconds) in TPS-based anomaly.

escalation-period

Specifies the escalation period (in seconds) in TPS-based anomaly.

geo-captcha-challenge

Enables or disables Geolocation-based CAPTCHA challenge in TPS-based anomaly.

geo-client-side-defense

Enables or disables Geolocation-based client side integrity defense in TPS-based anomaly.

geo-minimum-share

Specifies the minimum traffic share for detection in Geolocation detection criteria of TPS-based anomaly.

geo-rate-limiting

Enables or disables Geolocation-based rate limiting in TPS-based anomaly.

geo-request-blocking-mode

Specifies a Geolocation-based request blocking mode of TPS-based anomaly. The options are:

block-all

Specifies that the system blocks all requests from the respective Geolocation.

rate-limit

Specifies that the system blocks requests from the respective Geolocation based on the traffic share ratio. This is the default value.

geo-share-increase-rate

Specifies the percentage by which TPS increased in Geolocation detection criteria of TPS-based anomaly.

ip-captcha-challenge

Enables or disables Source IP-based CAPTCHA challenge in TPS-based anomaly.

ip-client-side-defense

Enables or disables Source IP-based client side integrity defense in TPS-based anomaly.

ip-maximum-tps

Specifies the amount which TPS reached in IP detection criteria of TPS-based anomaly.

ip-minimum-tps

Specifies the minimum TPS threshold for detection in IP detection criteria of TPS-based anomaly.

ip-rate-limiting

Enables or disables Source IP-based rate limiting in TPS-based anomaly.

ip-request-blocking-mode

Specifies a Source IP-based request blocking mode of TPS-based anomaly. The options are:

block-all

Specifies that the system blocks all requests from the respective Source IP address.

rate-limit

Specifies that the system blocks requests from the respective Source IP address based on the traffic share ratio. This is the default value.

ip-tps-increase-rate

Specifies the percentage by which TPS increased in IP detection criteria of TPS-based anomaly.

mode Specifies an operation mode of TPS-based anomaly. The options are:

- off Specifies that the system does not check for DoS attacks. This is the default value.

- transparent

Specifies that when the system detects an attack, it displays the attack data on the Reporting DoS Attacks screen. In transparent mode the system does not drop requests either from the attacking IP address, or to attacked URLs.

- blocking

Specifies that when the system detects an attack, in addition to displaying the attack data on the Reporting DoS Attacks screen, the system also drops either connections from the attacking IP address, or requests to attacked URLs.

site-captcha-challenge

Enables or disables Site-wide CAPTCHA challenge in TPS-based anomaly.

site-client-side-defense

Enables or disables Site-wide client side integrity defense in TPS-based anomaly.

site-maximum-tps

Specifies the amount which TPS reached in Site-wide detection criteria of TPS-based anomaly.

site-minimum-tps

Specifies the minimum TPS threshold for detection in Site-wide detection criteria of TPS-based anomaly.

site-rate-limiting

Enables or disables Site-wide rate limiting in TPS-based anomaly.

site-tps-increase-rate

Specifies the percentage by which TPS increased in Site-wide detection criteria of TPS-based anomaly.

static-url-mitigation

Enables or disables Static URL mitigation in TPS-based anomaly.

url-captcha-challenge

Enables or disables URL-based CAPTCHA challenge in TPS-based anomaly.

url-client-side-defense

Enables or disables URL-based client side integrity defense in TPS-based anomaly.

url-maximum-tps

Specifies the amount which TPS reached in URL detection criteria of TPS-based anomaly.

url-minimum-tps

Specifies the minimum TPS threshold for detection in URL detection criteria of TPS-based anomaly.

url-rate-limiting

Enables or disables URL-based rate limiting in TPS-based anomaly.

url-tps-increase-rate

Specifies the percentage by which TPS increased in URL detection criteria of TPS-based anomaly.

- trigger-irule

Specifies, when enabled, that the system activates an Application DoS iRule event. The default value is disabled.

- single-page-application

Specifies, when enabled, that the system supports a Single Page Applications. The default value is disabled.

- fastl4-acceleration-profile

Specifies a fastL4 profile that used for DOS acceleration. None - if disable acceleration.

- scrubbing-enable

Specifies whether to enable Traffic Scrubbing during attacks by advertising BGP routes. This requires configuration of security scrubber profile, and will function even when the mode is set to transparent.

- scrubbing-duration-sec

Specifies the duration of the Traffic Scrubbing BGP route advertisement, in seconds. This is used when scrubbing-enable is enabled.

- rtbh-enable

Specifies whether to enable Remote Triggered Black Hole (RTBH) of attacking IPs by advertising BGP routes. This requires configuration of security blacklist-publisher, and will function even when the Operation Mode is set to transparent.

- rtbh-duration-sec

Specifies the duration of the RTBH BGP route advertisement, in seconds. This is used when rtbh-enable is enabled.

- description

User defined description.

protocol-dns

Adds, deletes, or replaces a single Protocol DNS Security sub-profile. You can configure the following options for Protocol DNS Security:

name Specifies a dummy name for enabled Protocol DNS Security. This option is required for the operations create, delete, modify, and replace-all-with.

dynamic-signatures

Specifies options related to DNS Behavioral DoS (Dynamic Signatures) feature per virtual server by virtue of attaching a dos profile to a virtual server. Following options are configurable for this feature:

detection

Specifies the mode for detection of anomalies in traffic for the purpose of dynamic signature generation. Following modes are supported: disabled, enabled and learn-only.

Mode learn-only is same as enabled except that the system does not generate any logs (or alerts the user). It is used mainly to learn the baseline thresholds for the traffic.

Default is disabled.

mitigation

Specifies the mode for mitigation of anomalous traffic (specified in form of dynamic signatures). Following modes are supported: none, low, medium and high.

Each mode represents the severity (or aggressiveness) at which the system should try to mitigate the anomalous traffic.

Default is none.

multiplier-mitigation-percentage

Specifies the mitigation multiplier value of all the vectors in the dns dos profile in percentage when using manual-multiplier-mitigation mode.

dns-query-vector

Adds, deletes, or replaces Protocol DNS DoS vectors. You can configure the following options for DNS query vectors:

query-type

Specifies the vector (DNS query) type for DoS attack detection.

enforce

This option is deprecated in version 13.1.0 and is replaced by state. Enable or disable the packet drop action of DOS detection for this attack type.

auto-threshold

This option is deprecated in version 13.1.0 and is replaced by threshold-mode. Enables the auto threshold mode for dos detection and dos mitigation. The default value is disabled.

allow-upstream-scrubbing

Enables allow upstream scrubbing. The default value is disabled.

attacked-dst

Enables attacked-destination. The default value is disabled.

auto-scrubbing

Enables specifying destination IP scrubbing. The default value is disabled.

bad-actor

Enables per-source IP based bad actor detection

multiplier-mitigation-percentage

Specifies the mitigation multiplier value of this specific vector in percentage when using manual-multiplier-mitigation mode, The default value used is inherited from the dns dos profile.

per-source-ip-detection-pps

Bad actor detection rate (for single IP address) of this vector

per-source-ip-limit-pps

Bad actor allowed rate (for single IP address) of this vector

per-dst-ip-detection-pps

Specifies the attack detection threshold (pps) per destination IP. The default value is infinite.

per-dst-ip-limit-pps.

Specifies the attack mitigation threshold (pps) per destination IP. The default value is infinite.

scrubbing-category

Specifies per-DstIP scrubbing category. The default value is none.

scrubbing-detection-seconds

Specifies duration in seconds for which the destination IP has been offended/attacked. The default value is 10.

scrubbing-duration

Specifies duration in seconds for which this IP should be scrubbed. The default value is 900.

`rate-increase`

Specifies the rate increase for DoS attack detection.

`rate-limit`

Specifies the rate limit for DoS attack detection. If the value is infinite the detection is disabled.

`rate-threshold`

Specifies the rate threshold for DoS attack detection. If the value is infinite the detection is disabled.

`state`

Specifies the run time state of this signature. The options are the same as those in `network-attack-vector`.

`suspicious`

Specifies if the vector considers all packets or only unsolicited packets. The default value is false.

`threshold-mode`

Enables the threshold mode for dos detection and dos mitigation. The default value is manual. The options are the same as those in `network-attack-vector`.

`prot-err-attack-detection`

Specifies if protocol errors attack detection is enabled or not. Eg: Malformed, Malicious DoS attacks.

`prot-err-atck-rate-incr`

Specifies the protocol errors rate increase for DoS attack detection.

`protocol-sip`

Adds, deletes, or replaces a single Protocol SIP Security sub-profile. You can configure the following options for Protocol SIP Security:

`name` Specifies a dummy name for enabled Protocol SIP Security. This option is required for the operations create, delete, modify, and replace-all-with.

`prot-err-atck-rate-increase`

Specifies the protocol errors rate increase for DoS attack detection.

`prot-err-atck-rate-threshold`

Specifies the protocol errors rate threshold for DoS attack detection.

`prot-err-attack-detection`

Specifies if protocol errors attack detection is enabled or not. Eg: Malformed packets DoS attacks.

`multiplier-mitigation-percentage`

Specifies the mitigation multiplier value of all the vectors in the sip dos profile in percentage when using manual-multiplier-mitigation mode.

`sip-attack-vector`

Adds, deletes, or replaces Protocol SIP DoS vectors. You can configure the following options for SIP method vectors:

`type` Specifies the vector type (SIP method) for DoS attack detection.

`enforce`

This option is deprecated in version 13.1.0 and is replaced by `state`. Enable or disable the packet drop action of DOS detection for this attack type.

`auto-threshold`

This option is deprecated in version 13.1.0 and is replaced by `threshold-mode`. Enables the auto threshold mode for dos detection and dos mitigation. The default value is disabled.

`allow-upstream-scrubbing`

Enables allow upstream scrubbing. The default value is disabled.

`attacked-dst`

Enables attacked-destination. The default value is disabled.

`auto-scrubbing`

Enables specifying destination IP scrubbing. The default value is disabled.

`bad-actor`

Enables per-source IP based bad actor detection

`multiplier-mitigation-percentage`

Specifies the mitigation multiplier value of this specific vector in percentage when using manual-multiplier-mitigation mode, The default value used is inherited from the sip dos profile.

`per-source-ip-detection-pps`

Bad actor detection rate (for single IP address) of this vector

`per-source-ip-limit-pps`

Bad actor allowed rate (for single IP address) of this vector

`per-dst-ip-detection-pps`

Specifies the attack detection threshold (pps) per destination IP. The default value is infinite.

`per-dst-ip-limit-pps`
Specifies the attack mitigation threshold (pps) per destination IP. The default value is infinite.

`scrubbing-category`
Specifies per-DstIP scrubbing category. The default value is none.

`scrubbing-detection-seconds`
Specifies duration in seconds for which the destination IP has been offended/attacked. The default value is 10.

`scrubbing-duration`
Specifies duration in seconds for which this IP should be scrubbed. The default value is 900.

`rate-increase`
Specifies the rate increase for DoS attack detection.

`rate-limit`
Specifies the rate limit for DoS attack detection. If the value is infinite the detection is disabled.

`rate-threshold`
Specifies the rate threshold for DoS attack detection. If the value is infinite the detection is disabled.

`state`
Specifies the run time state of this signature. The options are the same as those in `network-attack-vector`.

`suspicious`
Specifies if the vector considers all packets or only unsolicited packets. The default value is false.

`threshold-mode`
Enables the threshold mode for dos detection and dos mitigation. The default value is manual. The options are the same as that in `network-attack-vector`.

`dos-network`
Adds, deletes, or replaces a single Network DoS Security sub-profile. You can configure the following options for Network DoS Security:

`name` Specifies a dummy name for enabled Network DoS Security. This option is required for the operations create, delete, modify, and replace-all-with.

`dynamic-signatures`
Specifies options related to L4 Behavioral DoS (Dynamic Signatures) feature per virtual server by virtue of attaching a dos profile to a virtual server. Following options are configurable for this feature:

`detection`
Specifies the mode for detection of anomalies in traffic for the purpose of dynamic signature generation. Following modes are supported: disabled, enabled and learn-only.

Mode learn-only is same as enabled except that the system does not generate any logs (or alerts the user). It is used mainly to learn the baseline thresholds for the traffic.

Default is disabled.

`mitigation`
Specifies the mode for mitigation of anomalous traffic (specified in form of dynamic signatures). Following modes are supported: none, low, medium and high.

Each mode represents the severity (or aggressiveness) at which the system should try to mitigate the anomalous traffic.

Default is none.

`scrubber-enable`
Specifies the configuration mode for enabling or disabling the feature to scrub the attack traffic upon dynamic signature match. Default is no.

`scrubber-category`
Specifies the IP Intelligence category used for scrubbing the attack traffic upon dynamic signature match that constitutes destination IP address component. Default category is `attacked_ips`.

`scrubber-advertisement-period`
Specifies the advertisement period for which the attack traffic is scrubbed. Default is 300 seconds.

`multiplier-mitigation-percentage`
Specifies the mitigation multiplier value of all the vectors in the network dos profile in percentage when using manual-multiplier-mitigation mode.

`network-attack-vector`
Adds, deletes, or replaces Network Attack DoS vectors. You can configure the following options for Network Attack vectors:

attack-type
Specifies the vector type (Network Attack) for DoS attack detection.

enforce
This option is deprecated in version 13.1.0 and is replaced by state. Enable or disable the packet drop action of DOS detection for this attack type.

auto-threshold
This option is deprecated in version 13.1.0 and is replaced by threshold-mode. Enables the auto threshold mode for dos detection and dos mitigation. The default value is disabled.

rate-increase
Specifies the rate increase for DoS attack detection.

rate-limit
Specifies the rate limit for DoS attack detection. If the value is infinite the detection is disabled.

rate-threshold
Specifies the rate threshold for DoS attack detection. If the value is infinite the detection is disabled.

packet-types
Specifies the packet types for Sweep attack vector.

allow-upstream-scrubbing
Enables allow upstream scrubbing. The default value is disabled.

attacked-dst
Enables attacked-destination. The default value is disabled.

auto-scrubbing
Enables specifying destination IP scrubbing. The default value is disabled.

bad-actor
Enables per-source IP based bad actor detection

multiplier-mitigation-percentage
Specifies the mitigation multiplier value of this specific vector in percentage when using manual-multiplier-mitigation mode, The default value used is inherited from the network dos profile.

per-source-ip-detection-pps
Bad actor detection rate (for single IP address) of this vector

per-source-ip-limit-pps
Bad actor allowed rate (for single IP address) of this vector

per-dst-ip-detection-pps
Specifies the attack detection threshold (pps) per destination IP. The default value is infinite.

per-dst-ip-limit-pps
Specifies the attack mitigation threshold (pps) per destination IP. The default value is infinite.

scrubbing-category
Specifies per-DstIP scrubbing category. The default value is none.

scrubbing-detection-seconds
Specifies duration in seconds for which the destination IP has been offended/attacked. The default value is 10.

scrubbing-duration
Specifies duration in seconds for which this IP should be scrubbed. The default value is 900.

state
Specifies the run time state of this signature.

The options are:

- disabled**
Do not learn, do not collect stats.
- learn-only**
Learn/Collect stats, but do not "detect" ("alarm" in ASM-speak) any attacks,
- detect-only**
Learn/Collect stats/detect, but do not mitigate (rate-limit/drop, challenge, etc.) any attacks.
- mitigate**
Learn/Collect stats/detect/mitigate (using whichever mitigations are configured).

threshold-mode
Enables the threshold mode for dos detection and dos mitigation. The default value is manual.

The options are:

`manual`

Specifies the manual thresholds.

`stress-based-mitigation`

Specifies the manual detection ("alarm") threshold, but mitigation threshold is stress-based.

`fully-automatic`

Specifies both the detection ("alarm") and mitigation thresholds are automatically computed.

`manual-multiplier-mitigation`

Specifies the detection ("alarm") threshold is automatically computed. The mitigation threshold is calculated by the detection threshold multiplies the multiplier-mitigation-percentage.

`whitelist`

Specifies the Dos srcIP whitelist configuration.

`http-whitelist`

Specifies the IP addresses and subnets whitelist configuration for Application Security (Overrides the global whitelist).

`custom-signatures`

Specifies options related to L4 Behavioral DoS Signatures feature per virtual server by virtue of attaching one or more signatures objects. Following options are configurable for this feature:

`threshold-mode`

Specifies the mode for setting the rate limit thresholds to be used for the matching traffic. Following modes are supported: manual, fully-automatic and stress-based-mitigation. Default is manual.

`state`

Specifies the operational state of the attached signature. The states supported are: disabled, learn-only, detect-only and mitigate. Default is disabled.

`suspicious`

Specifies if the vector considers all packets or only unsolicited packets. The default value is false.

`manual-detection-threshold`

Specifies the attack detection threshold of the attached signature.

Default is infinite.

`manual-mitigation-threshold`

Specifies the attack mitigation threshold of the attached signature.

Default is infinite.

`glob` Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

`name` Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

`partition`

Displays the administrative partition within which the component resides.

`regex`

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, edit, glob, list, ltm virtual, modify, regex, security, security dos, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2019-09-08 security dos profile(1)

NAME

spva-stats - Shows the Configuration and Data path stats for sPVA.

MODULE

security dos

SYNTAX

Displays the sPVA stats

DISPLAY

show security dos spva-stats

DESCRIPTION

This module shows the sPVA related configuration and data path stats. sPVA stats are relevant only in case the Hardware is sPVA capable.

EXAMPLES

show security dos spva-stats

Displays the below sPVA reslated stats

Total SW entry

Total sPVA entry count present in SW

Total HSB entry

Max number of sPVA entries in HSB

Used HSB entry

Total number of sPVA entries present in HSB

Global white list SW entry

Global whitelist entries present in SW

Global white list HSB entry

Global Whitelist entries present in HSB.

Global white list total entry

Total number of Global white list entries present in SW and HSB

Global black list SW entry

Total number of global back list sPVA entries present in SW

Global black list HSB entry

Total number of global black list sPVA entries present in HSB

Global black list total entry

Total number of global black list entry count present in SW and HSB

VS white list SW entry

Total whitelist SW entry count for virtual server with port as wildcard.

VS white list HSB entry

Total whitelist HSB entry count for virtual server with port as wildcard.

VS white list total entry

Total count of white list entries present in SW and HSB for virtual server with port as wildcard.

VS black list SW entry

Total black list SW entry count for virtual server with port as wildcard.

VS black list HSB entry

Total black list HSB entry count for virtual server with port as wildcard.

VS black list total entry

Total count of black list entries present in SW and HSB for virtual server with port as wildcard.

App white list SW entry

Total whitelist SW entry count for application virtual server.

App white list HSB entry

Total whitelist HSB entry count for application virtual server.

App white list total entry

Total count of white list entries present in SW and HSB for application virtual server

App black list SW entry

Total black list SW entry count for application virtual server

App black list HSB entry

Total black list HSB entry count for application virtual server.

App black list total entry

Total count of black list entries present in SW and HSB for application virtual server

Global white list SW hit

Total packets hits for global white list in SW

Global white list HSB hit
Total packets hits for global white list in HSB

Global white list total hit
Total packets hits for global white list in SW and HSB.

Global black list SW hit
Total packets hits for global black list in SW

Global black list HSB hit
Total packets hits for global black list in HSB

Global black list total hit
Total packets hits for global black list in SW and HSB.

VS white list SW hit
Total packet hits for white list in SW for virtual server with port as wildcard.

VS white list HSB hit
Total packet hits for white list in HSB for virtual server with port as wildcard.

VS white list total hit
Total packet hits for white list in SW and HSB for virtual server with port as wildcard.

VS black list SW hit
Total packet hits for black list in SW for virtual server with port as wildcard.

VS black list HSB hit
Total packet hits for black list in HSB for virtual server with port as wildcard.

VS black list total hit
Total packet hits for black list in SW and HSB for virtual server with port as wildcard.

App white list SW hit
Total packet hits for white list in SW for application virtual server.

App white list HSB hit
Total packet hits for white list in HSB for application virtual server.

App white list total hit
Total packet hits for white list in SW and HSB for application virtual server.

App black list SW hit
Total packet hits for black list in SW for application virtual server.

App black list HSB hit
Total packet hits for black list in HSB for application virtual server.

App black list total hit
Total packet hits for black list in SW and HSB for application virtual server.

White list SW hit
Total packet hits for white list in SW

White list HSB hit
Total packet hits for white list in HSB

White list total hit
Total packet hits for white list in SW and HSB

Black list SW hit
Total packet hits for black list in SW

Black list HSB hit
Total packet hits for black list in HSB

Black list total hit
Total packet hits for black list in SW and HSB

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2015. All rights reserved.

BIG-IP 2016-01-07 security dos spva-stats(1)

NAME

stress-stats - Displays DoS stress statistics, or runs DoS stress statistics to set stress calculation mode on the BIG-IP(r) system.

MODULE

security dos

SYNTAX

Display and run DoS stress statistics using the following syntax.

MODIFY

run stress-stats

options:

context-name [string]

stress [0-100 | auto]

DISPLAY

show stress-stats

options:

field-fmt

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the stress-stats component to display DoS stress-stats statistics, or run dos stress-stats command to set auto or manual mode on stress calculation.

EXAMPLES

```
run stress-stats context-name vs1 stress auto
```

Makes the stress level calculation to be in auto-mode on virtual server/context vs1 that have DoS Profile attached.

```
run stress-stats context-name vs1 stress 10
```

Makes the stress level calculation to be in manual-mode with value 10% on virtual server/context vs1 that have DoS Profile attached.

```
show stress-stats
```

Shows the stress-stats statistics on all virtual servers.

```
show stress-stats field-fmt
```

Shows the stress-stats statistics on all virtual servers in field format.

OPTIONS

Use these options to control stress-stats of the security dos:

context-name

Specifies the unique name for the context in run command. This option is required.

stress

Specifies the stress value for the context in run command. This option is required. The range is from 0 to 100 (0% to 100%) for manual-mode or auto for auto-mode.

SEE ALSO

show, run, security dos stress-stats

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2019. All rights reserved.

BIG-IP 2019-06-20 security dos stress-stats(1)

security dos udp-portlist

NAME

udp-portlist - Configures the DoS udp portlist component within the security dos module using the syntax shown in the following sections. These DoS udp portlist entries are applied to all udp packets except those going through the management interface.

MODULE

security dos

SYNTAX

MODIFY

```

modify udp-portlist dos-udp-portlist
options:
description [string]
list-type [exclude-listed-ports | include-listed-ports]
entries [modify | replace-all-with] {
  [entry] {
    options:
description [string]
match-direction [both | dst | none | src]
port-number [number]
  }
}

```

```

DISPLAY
list udp-portlist

```

DESCRIPTION

You can use the `udp-portlist` component to configure a DoS UDP portlist of upto eight entries for all UDP traffic except the management interface. The HSB hardware compares all incoming UDP traffic to the `udp-portlist` entries. There are 2 types of behavior, depending upon whether the `udp` port list is configured as a white list or as a black list. White list and black list are mutually exclusive properties of a UDP port list.

If the `udp` port list is configured as a `list-type` of `exclude-listed-ports`, and if a match is found on an incoming packet, then we do not increment the UDP Flood DoS vector. If a match is not found, then the UDP Flood DoS vector checks are done on those packets.

If the `udp` port list is configured as a `list-type` of `include-listed-ports`, and if a match is found on an incoming packet, then we increment the UDP Flood DoS vector. If a match is not found, then the UDP Flood DoS vector checks are not done on the packets.

Either destination port or source port or both can be specified in a `udp-portlist` entry.

EXAMPLES

```

modify udp-portlist dos-udp-portlist description "bad ports" list-type include-listed-ports
Modifies the udp-portlist dos-udp-portlist to a blacklist.

```

```

modify udp-portlist dos-udp-portlist list-type exclude-listed-ports
Modifies the udp-portlist dos-udp-portlist to a white-list.

```

```

modify udp-portlist dos-udp-portlist description "bad ports" entries modify { entry1 { match-direction src
port-number 161 } }
Modifies an entry. The new entry is for source UDP port 161. It matches any UDP packet whose source port is 161.

```

```

modify udp-portlist dos-udp-portlist entries modify { 161 { match-direction both } }

```

Modifies the entry for destination UDP port 161 to source and destination port 161. It matches any UDP packet whose destination or source port is 161.

```

security dos udp-portlist dos-udp-portlist {
entries {
entry1 {
match-direction both
port-number snmp
}
entry2 { }
entry3 { }
entry4 { }
}
white-list
}

```

Displays the current list of DoS UDP portlist entries.

OPTIONS

`description`
Your description for the DoS `udp-portlist`.

`list-type`
Sets the list type to be either `exclude-listed-ports` or `include-listed-ports`

`include-listed-ports`
Sets the property of the `dos-udp-portlist` list to `include-listed-ports` (Blacklist).

`exclude-listed-ports`
Sets the property of the `dos-udp-portlist` list to `exclude-listed-ports` (Whitelist).

`entries`
Modifies a `udp-portlist` entry.

`modify`
Modifies the existing entry that you specify next, in curly braces (`{}`). After the entry name, enter the new configuration (port mode and port number) settings for the entry inside a nested set of curly braces.

`replace-all-with`
Replaces the current set of `udp-portlist` entries with the entry(s) that you specify next, in curly braces (`{}`).

Enter the name of a entry to be modified, then enter an open curly brace ({), one or more of the following options, and a closed curly brace (}).

`description`

Your description for the current entry.

`match-direction`

Set the mode of matching (source, destination or both).

`port-number`

Set the port number for matching.

SEE ALSO

`edit`, `list`, `modify`, `security`, `security dos`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

BIG-IP 2016-03-14 security dos udp-portlist(1)

security dos virtual

NAME

virtual - Displays and resets dos per-virtual statistics, or run dos per-virtual command on the BIG-IP(r) system.

MODULE

security dos

SYNTAX

Manage the virtual component within the security dos module using the syntax in the following sections.

DISPLAY

`show virtual name [vs_name]`

options:

`dns-nxdomain-stat`

`field-fmt`

`query-valid-domain [fqdn]`

MODIFY

`reset-stats virtual name [vs_name]`

options:

`dns-nxdomain-stat`

RUN

`run virtual name [vs_name]`

options:

`auto-threshold-relearn`

`dns-nxdomain-relearn`

`dynamic-signatures-history-relearn`

DESCRIPTION

You can use the virtual component to display and reset dos per-virtual statistics, or run dos per-virtual commands.

EXAMPLES

`show virtual name my_virtual query-valid-domain www.f5.com`

Displays if a FQDN(Fully Qualified Domain Name) is present in the per-virtual learned valid domains.

`show virtual name my_virtual dns-nxdomain-stat`

Displays the dns-nxdomain statistics on the specified virtual server.

`reset-stats virtual name my_virtual dns-nxdomain-stat`

Resets the dns-nxdomain-stat statistics on the specified virtual server.

`run virtual name my_virtual auto-threshold-relearn`

Clears the auto-threshold history on the specified virtual server auto-threshold vector.

`run virtual name my_virtual dns-nxdomain-relearn`

Clears the dns-nxdomain history on the specified virtual server.

```
run virtual dns-nxdomain-relearn
```

Clears the dns-nxdomain history on all virtual servers.

```
run virtual name my_virtual dynamic-signatures-history-relearn
```

Clears the dynamic-signatures history on the specified virtual server dynamic-signatures vector.

OPTIONS

name Apply the command to the specified virtual server. It can be optional in some commands/options. In these cases, the command will be applied to all virtual servers.

show command: options are

dns-nxdomain-stat: name is required.

field-fmt: only work with dns-nxdomain-stat.

query-valid-domain [fqdn]: name is required.

run command: options are

auto-threshold-relearn: name is required

dns-nxdomain-relearn: name is optional. The command applies to all virtual servers if name is not specified.

auth-threshold-relearn: name is required.

reset-stats command: options are

dns-nxdomain-stat: name is required.

SEE ALSO

reset-stats, run, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2017. All rights reserved.

BIG-IP 2018-01-24 security dos virtual(1)

security firewall address-list

NAME

address-list - Configures an address-list for use by firewall rules. An address list is a list of IP-address prefixes to compare against the source-IP address and/or destination-IP address in an IP packet.

MODULE

security firewall

SYNTAX

CREATE/MODIFY

```
create address-list [name]
```

```
modify address-list [[name] | all]
```

options:

```
addresses [add | delete | modify | replace-all-with] {  
  [ [ip address] ]  
}
```

```
fqdns [add | delete | replace-all-with] {  
  [ fully qualified domain names]  
}
```

```
fqdns none
```

```
geo [add | default | delete | replace-all-with] {  
  [ [country_code[:state_name/city_name] ] ]  
}
```

```
geo none
```

```
app-service [name]
```

```
description [string]
```

```
edit address-list [[name] | all]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list address-list [[name] | all | [property]]
```

```
show running-config address-list [[name] | all | [property]]
```

DELETE

```
delete address-list [[name] | all]
```

DESCRIPTION

You can use the address-list component to define reusable lists of addresses. You can use an address list in any of the following firewalls and firewall rule lists: net self, net route-domain, security firewall global-rules, security firewall rule-list, security firewall management-ip-rules, and ltm virtual. A firewall rule compares all of the addresses in the list to either the source or destination IP in the packet, depending on how you apply the list. If there is a match, the firewall rule takes an action, such as accepting or dropping the packet.

EXAMPLES

```
create address-list alist1 addresses add { 10.10.1.1 10.10.1.2 192.168.24.0/24 }
```

Creates a new address list, "alist1," with two IPv4 addresses and one IPv4 subnet.

```
modify address-list alist1 addresses modify { 10.10.1.1 { description "management IP at wwmed site3" } }
```

Modifies the above address list with a description for the first address.

```
modify alist1 geo add { TR:Istanbul }
```

Modifies the above address list with an addition of a country:city/state.

```
modify address-list alist1 addresses add { 2001:DB8:a::/64 }
```

Modifies the same address list by adding an IPv6 subnet.

```
list address-list alist1
security firewall address-list alist1 {
  addresses {
    10.10.1.1 {
      description "management IP at wwmed site3"
    }
    10.10.1.2 { }
    192.168.24.0/24 { }
    2001:db8:a::/64 { }
  }
}
```

Shows the modified address list.

```
create address-list xyz fqdns add { xyz.com }
```

Creates a new address list, "xyz" with a single fully qualified domain 'xyz.com'.

```
modify address-list xyz addresses add { 2001:DB8:a::/64 } fqdns add { abc.com }
```

Modifies the same address list by adding an IPv6 subnet and another fully qualified domain 'abc.com'.

```
list address-list xyz
security firewall address-list xyz {
  addresses {
    2001:db8:a::/64 { }
  }
  fqdns {
    abc.com { }
    xyz.com { }
  }
}
```

Shows the above address list 'xyz'.

OPTIONS

addresses

Specifies a list of IP addresses and/or subnets to compare against a packet's source or destination address. The format for an IPv4 address is a.b.c.d[/prefix]. The general format for an IPv6 address is a:b:c:d:e:f:g:h[/prefix]; you can shorten this by eliminating leading zeros from each field (for example, you can shorten "2001:0db7:3f4a:09dd:ca90:ff00:0042:8329" to "2001:db7:3f4a:9dd:ca90:ff00:42:8329"), and/or by removing the longest contiguous field of zeros (for example, you can shorten "2001:0:0:0:c34a:0:23ff:678" to "2001::c34a:0:23ff:678"). TMSH accepts any valid text representation of IPv6 addresses, as defined in RFC 2373 (see).

The next keyword specifies the action to take with the addresses (add, delete, modify, or replace the current set of addresses).

add Creates a new address list, which you specify next with IP addresses and/or prefixes in curly braces ({}).

delete
Deletes the address(es) that you specify next, in curly braces ({}).

modify
Makes it possible to replace the optional description(s) for the address(es). You can specify a description in a nested set of curly braces after each address.

replace-all-with
Replaces the current set of IP addresses with the address(es) that you specify next, in curly braces ({}).

fqdns

Specifies a list of fully qualified domain names to compare against packet's destination IP address domain.

The next keyword specifies the action to take with the fqdns (add, delete, or replace the current set of fqdns).

`geo` Specifies a list of geographic locations that the packet will be compared against.

`app-service`

Associates this address list with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers. The asm module (see asm) has components for working with iApps.

`description`

Is your description for this address list.

SEE ALSO

edit, list, modify, net self, net route-domain, security firewall global-rules, security firewall management-ip-rules, security firewall rule-list, ltm virtual, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 security firewall address-list(1)

security firewall config-change-log

NAME

config-change-log - Configures firewall configuration change log setting.

MODULE

security firewall

SYNTAX

Modify the config-change-log component within the security firewall module using the syntax shown in the following sections.

MODIFY

modify config-change-log log-changes [automatic | on | off]

modify config-change-log log-publisher [none | [name]]

DISPLAY

list config-change-log

show running-config config-change-log

DESCRIPTION

You can use the config-change-log component to configure if changes to the firewall rules should be logged or not. The default is to automatically determine if log is needed based on the mode of the on-demand-compilation. If the mode of the on-demand-compilation is 'enabled', the changes will be logged. If change log is enabled, a publisher need to be configured too.

EXAMPLES

modify config-change-log log-changes on log-publisher local-db-publisher

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-09-16 security firewall config-change-log(1)

security firewall container-stat

NAME

container-stat - Show the compilation result of firewall rules.

MODULE
security firewall

SYNTAX
show container-stat

DESCRIPTION
You can use the container-stat component to display the compilation result of firewall rules for each container. For firewall rules, a container is uniquely identified by the combination of context type, context name and policy type. The outputs of the command include: the time it takes to compile the rules and to perform overlapping check for each container, the number of micro-rules, the size of the containers in bytes, the amount of memory used for compilation, and the time the container is activated.

EXAMPLES
show security firewall container-stat

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015. All rights reserved.

BIG-IP 2015-03-11 security firewall container-stat(1)

security firewall context-stat

NAME
context-stat - Displays and resets firewall statistics of the specified context on the BIG-IP(r) system. You can only use the show and reset-stats command with this component.

MODULE
security firewall

SYNTAX
show context-stat
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt
reset-stats context-stat

DESCRIPTION
You can use the context-stat component to display or reset firewall statistics of the specified context.

EXAMPLES
show context-stat
Displays firewall rule's statistics in the system default units.
show context-stat raw
Displays raw firewall rule's statistics.
reset-stats context-stat
Resets firewall statistics for all contexts.

SEE ALSO
show, reset-stats, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2017. All rights reserved.

BIG-IP 2017-09-06 security firewall context-stat(1)

security firewall current-state

NAME

current-state - Show the current state of firewall rules compilation.

MODULE

security firewall

SYNTAX

show current-state

DESCRIPTION

You can use the current-state component to display the current system-wide state of firewall rules compilation. The outputs of the command includes: the mode of on-demand compilation and on-demand rule deploy (enabled or disabled), the rule compiler status (quiescent, pending-compilation, being-compiled, pending-deployment, being-deployed, failed-compilation, failed-deployment, pccd-failed), compilation start time, compilation end time, deployment start time, deployment end time, aggregate number of micro-rules, active blob name, blob creation time, and if the blob MD5 is verified or not.

EXAMPLES

show security firewall current-state

SEE ALSO

security firewall on-demand-compilation, security firewall on-demand-rule-deploy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015. All rights reserved.

BIG-IP 2015-04-27 security firewall current-state(1)

security firewall fqdn-entity

NAME

fqdn-entity - Perform on-demand refresh to query IP mappings for one (or all) configured FQDNs in firewall rules.

MODULE

security firewall

SYNTAX

Use the fqdn-entity component within the security firewall module to perform on-demand refresh to query IP mappings for a specific FQDN (or all FQDNs) configured in firewall rules using the following syntax:

DISPLAY

load fqdn-entity [all | name]

DESCRIPTION

You can use the fqdn-entity component to perform on-demand refresh to query IP mappings for one or all configured FQDNs in firewall rules.

BigIP will periodically refresh IP mappings for each FQDN upon previous mappings expiry time or after refresh-interval as configured in security firewall global-fqdn-policy whichever is greater. fqdn-entity component can be used to refresh the IP mappings irrespective of either the expiry time or the refresh-interval.

EXAMPLES

load fqdn-entity all

Perform on-demand refresh for all the configured FQDNs in firewall rules.

load fqdn-entity f5.com

Perform on-demand refresh for fqdn 'f5.com' configured in one of the firewall rules.

SEE ALSO

security firewall address-list, security firewall policy, security firewall rule-list, security firewall fqdn-info, security firewall global-fqdn-policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2015. All rights reserved.

BIG-IP 2015-07-22 security firewall fqdn-entity(1)

security firewall fqdn-info

NAME

fqdn-info - Query run time information for one (or all) configured FQDNs in firewall rules.

MODULE

security firewall

SYNTAX

Use the fqdn-info component within the security firewall module to query run time information for one (or all) configured FQDNs in firewall rules using the following syntax:

DISPLAY

```
show fqdn-info fqdn [all | name]
```

DESCRIPTION

You can use the fqdn-info component to query run time information for one or all configured FQDNs in firewall rules. Use option fqdn to query information regarding a specific FQDN or all FQDNs configured in firewall rules.

Following information related to each FQDN is queried using fqdn-info component:

- a) Last Refresh time.
- b) Next Refresh time.
- c) Last Successful Refresh.
- d) DNS Records valid until
- e) IP Addresses

OPTIONS

fqdn Specifies the name of a FQDN. Use 'all' to query information for all configured FQDNs.

EXAMPLES

```
show fqdn-info fqdn all
```

Query information for all the configured FQDNs in firewall rules.

```
show fqdn-info fqdn f5.com
```

Query information for fqdn 'f5.com' configured in one of the firewall rules.

SEE ALSO

security firewall address-list, security firewall policy, security firewall rule-list

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2014. All rights reserved.

BIG-IP 2014-10-08 security firewall fqdn-info(1)

security firewall global-fqdn-policy

NAME

global-fqdn-policy - Configures the global fqdn policy which is used to resolve FQDN names to IP Address mappings for the FQDN names that are specified in the firewall rules.

MODULE

security firewall

SYNTAX

Modify the global-fqdn-policy component within the security firewall module using the syntax shown in the following sections.

MODIFY

```
modify global-fqdn-policy
```

options:

```
app-service [name]
```

```
description [string]
```

```
dns-resolver [ [resolver_name] | none ]
```

```
refresh-interval [integer]
```

```
edit global-fqdn-policy
```

options:

all-properties
non-default-properties
one-line
partition
recursive

DISPLAY
list global-fqdn-policy
show running-config global-fqdn-policy
options:
all-properties
non-default-properties
one-line
partition
recursive

DESCRIPTION

You can use the global-fqdn-policy component to configure a net dns-resolver that will be used by firewall to resolve FQDN names to IP Address mappings. These mappings in turn will be used to match firewall rules (across all policies on all contexts) based on FQDN constraints.

EXAMPLES

```
modify global-fqdn-policy dns-resolver xyz
```

Modifies the global-fqdn-policy to use dns resolver object named 'xyz'. Default refresh-interval is 60 seconds.

```
modify global-fqdn-policy dns-resolver xyz refresh-interval 120
```

Modifies the global-fqdn-policy to use dns resolver object named 'xyz' and specify periodic refresh rate of 120 seconds (2 minutes) to re-resolve FQDN-to-IP mappings.

```
list global-fqdn-policy
```

Displays the current list of global-fqdn-policy contents.

OPTIONS

app-service
Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description
User defined description.

dns-resolver
Specifies an existing net dns-resolver. This will be used by firewall to obtain FQDN-to-IP Address mappings which will be used to match firewall rules based on FQDN constraints. Note dns-resolver none can be used to remove the object from global-fqdn-policy if and only if there are no AFM rules with (non empty) FQDN constraints.

refresh-interval
Specifies refresh interval to be used to re-resolve FQDN-to-IP mappings. Unit is in seconds and default is 60 seconds. Minimum allowed is 5 seconds and maximum is 2,764,800 (=32 days) seconds.

SEE ALSO

create, edit, list, modify, security firewall, security firewall policy, net dns-resolver tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2015. All rights reserved.

BIG-IP 2018-10-12 security firewall global-fqdn-policy(1)

security firewall global-rules

NAME

global-rules - Configures the global network firewall rules. These firewall rules are applied to all packets except those going through the management interface. They are applied first, before any firewall rules for the packet's virtual server, route domain, and/or self IP.

MODULE

security firewall

SYNTAX

MODIFY

```
modify global-rules
options:
  description [string]
  enforced-policy [ [policy_name] | none ]
  staged-policy [ [policy_name] | none ]
  service-policy [ [policy_name] | none ]
```

```
edit global-rules
options:
  all-properties
  non-default-properties
```

```
reset-stats global-rules
enforced-policy-rules { [rule name] }
staged-policy-rules { [rule name] }
```

```
options:
  fw-context-stat
  port-misuse
```

```
DISPLAY
list global-rules
show running-config global-rules
```

```
show global-rules
active
enforced-policy-rules
staged-policy-rules
```

```
options:
  fw-context-stat
  port-misuse
  overlapping-status
```

DESCRIPTION

You can use the global-rules component to configure network firewall policy which is enforced or staged on all IP and ICMP traffic except traffic on the management IP.

EXAMPLES

```
list global-rules
```

```
security firewall global-rules {
  enforced-policy /Common/policy1
}
```

Displays the current list of global rules.

OPTIONS

description

Your description for the global list of firewall rules.

enforced-policy

Specifies an enforced firewall policy. enforced-policy rules are enforced globally.

enforced-policy-rules

Specifies firewall rules enforced on traffic globally via referenced enforced-policy.

overlapping-status

Display detail overlapping information

port-misuse

Used to show or reset global port misuse policy statistics.

fw-context-stat

Used to show or reset firewall statistics for the global rules.

staged-policy

Specifies a staged firewall policy. staged-policy rules are not enforced while all the visibility aspects namely statistics, reporting and logging function as if the staged-policy rules were enforced globally.

staged-policy-rules

Specifies firewall rules staged on traffic globally via referenced staged-policy.

service-policy

Specifies a service policy that would apply to traffic globally. The service policy is applied to all flows, provided if there are no other context specific service policy configuration that overrides the global service policy. For example, when a service policy is configured both at a global level, as well as on a firewall rule, and a flow matches the rule, the more specific service policy configuration in the rule will override the service policy setting at the global level. The service policy associated here can be created using net service-policy command.

SEE ALSO

edit, list, modify, security firewall address-list, security firewall port-list, security firewall rule-list, security log profile, security firewall schedule, tmsh, security firewall policy, net service-policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2017-09-06 security firewall global-rules(1)

security firewall ipi-category-info

NAME ipi-category-info - Query firewall policy names and virtual-server names based on the given ipi-category, also query firewall rules based on ipi-category and firewall policy.

MODULE
security firewall

SYNTAX
Use the ipi-category-info component within the security firewall module to query firewall policies and virtual-servers with ipi category name and query firewall rules with ipi category name and the firewall policy with the following syntax:

DISPLAY
show security firewall ipi-category-info name [name of the ipi-category] [policies | virtuals] show security firewall ipi-category-info name [name of the ipi-category] rules [policy-name]

DESCRIPTION
You can use the ipi-category-info component to query firewall policies or virtual servers based on the given ipi-category. Also to query firewall rules based on the given ipi category and firewall policy.

OPTIONS
policy-name
Specifies the firewall policy-name.

EXAMPLES
show security firewall ipi-category-info name botnets policies
Returns firewall policies that are configured with 'botnets' ipi category
show security firewall ipi-category-info name botnets virtuals
Returns virtual-server names that are configured with 'botnets' ipi category
show security firewall ipi-category-info name botnets rules p1
Returns firewall rules that are under policy p1 with 'botnets' ipi category

SEE ALSO
security firewall fqdn-info, security firewall policy

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2019. All rights reserved.

BIG-IP 2019-04-22 security firewall ipi-category-info(1)

security firewall management-ip-rules

NAME
management-ip-rules - Configures the management IP firewall rules. These firewall rules are applied to all packets that go through the management interface.

MODULE
security firewall

SYNTAX
MODIFY
modify management-ip-rules
options:
description [string]
rules [add | delete | modify | replace-all-with] {
[[name]] }

```

options:
action [accept | accept-decisively | drop | reject]
description [string]
destination {
  address-lists [add | default | delete | replace-all-with] {
    [address list names...]
  }
  address-lists none
  addresses [add | default | delete | replace-all-with] {
    [ [ip address] | [ip address/prefixlen] ]
  }
  addresses none
  port-lists [add | default | delete | replace-all-with] {
    [port list names...]
  }
  port-lists none
  ports [add | default | delete | none | replace-all-with] {
    [ [port] | [port1-port2] ]
  }
  ports none
}
icmp [add | delete | modify | replace-all-with] {
  [ [icmp_type] | icmp_type:icmp_code ] {
    description [string]
  }
}
icmp none
ip-protocol [protocol name]
log [no | yes]
place-after [first | last | [rule name]]
place-before [first | last | [rule name]]
rule-list [rule list name]
schedule [schedule name]
source {
  address-lists [add | default | delete | replace-all-with] {
    [address list names...]
  }
  address-lists none
  addresses [add | default | delete | replace-all-with] {
    [ [ip address] | [ip_address/prefixlen] ]
  }
  addresses none
  port-lists [add | default | delete | replace-all-with] {
    [port list names...]
  }
  port-lists none
  ports [add | default | delete | replace-all-with] {
    [ [port] | [port1-port2] ]
  }
  ports none
  vlans [add | default | delete | replace-all-with] {
    [vlan names...]
  }
  vlans none
}
status [disabled | enabled | scheduled]
uuid [ | none | auto-generate]
}
}
rules none

```

edit management-ip-rules

```

options:
  all-properties
  non-default-properties

```

DISPLAY

list management-ip-rules

show running-config management-ip-rules

DESCRIPTION

You can use the management-ip-rules component to configure network firewall rules that are applied to all management interface traffic. The network software compares IP packets to the criteria specified in these rules. If a packet matches the criteria then the system takes the action specified by the rule. If a packet does not match a rule then the software compares the packet against the next rule. If a packet does not match any rule the packet is accepted.

For configuration sync management-ip-rules are synced to the devicegroup that has a type field of sync-failover. See "cm config-sync".

MATCHING AN IP PACKET

You can use this TMSH component to match against any or all of the following properties of an IP packet:

```

source address
source port
the packet's source VLAN
destination address

```

destination port
the higher-level protocol in the packet's payload

If you match against more than one of these items, a packet must pass all of your tests to successfully match. For example, if you match against a source subnet and several destination ports, a packet must originate from the given subnet and must also have one of the specified destination ports.

RULE ORDER

Rules are evaluated in the order that you specify. You can use the `list management-ip-rules` command to see the current rule order. As you add or modify rules in this component, you can use the `place-before rule-name` or `place-after rule-name` option to choose the rule's place in the sequence.

Rule order can determine whether or not a packet is dropped. Consider the following rules:

`rule_a`, matches source addresses against 172.16.0.0 and ACCEPTS all packets that match.
`rule_d`, matches source addresses against 172.16.39.0 and DROPS all packets that match.

Also consider a packet from a host at 172.16.39.55. If `rule_a` appears before `rule_d` in the rule list, the packet's source address matches `rule_a` first and the software accepts it. The software never reaches `rule_d` for comparison. If `rule_d` appears first instead, the packet's source address now matches `rule_d`; in this case, the software drops the packet.

EXAMPLES

```
modify management-ip-rules rules add { reject-internal-net { source { addresses replace-all-with { 172.27.0.0/16 } } } action reject place-before first } }
```

Creates a rule entry at the beginning of the list that rejects traffic from the 172.27.0.0 network.

```
modify management-ip-rules rules add { reject-insecure-ports { rule-list block_bad_mgmt place-before first } }
```

Adds a sub rule list to the management-IP firewall. Use the "security firewall rule-list" component to create a custom rule list.

```
list management-ip-rules
security firewall management-ip-rules {
  rules {
    reject-insecure-ports {
  rule-list block_bad_mgmt
    }
    reject-internal-net {
  action reject
  source {
    addresses {
  172.27.0.0/16 { }
    }
  }
}
}
```

Displays the current list of management-firewall rules.

```
modify management-ip-rules rules delete { reject-internal-net }
```

Removes the `reject-internal-net` rule from the management-IP firewall.

OPTIONS

`description`
Your description for the management-firewall rules.

`rules`
Adds, deletes, or replaces a firewall rule.

`add` Creates a new rule, which you specify next with a unique string in curly braces (`{}`). Use the `place-before` or `place-after` option inside the curly braces to determine the order of the rule. If this is the first rule, use the `replace-all-with` option instead of `add`.

`delete`
Deletes the rule that you specify next, in curly braces (`{}`).

`modify`
Modifies the existing rule that you specify next, in curly braces (`{}`). After the rule name, enter the new configuration settings for the rule inside a nested set of curly braces.

`replace-all-with`
Replaces the current set of global rules with the rule(s) that you specify next, in curly braces (`{}`). Use this option for the first management rule.

`none` Empties the list of management-firewall rules. This implicitly accepts all packets on the management interface.

Enter the name of a rule to be added or modified, then enter an open curly brace (`{`), one or more of the following options, and a closed curly brace (`}`).

`action`
Specifies the action that the system takes when a packet matches the rule.

accept

Specifies that a matching packet should be accepted. The security software stops comparing a matching packet to any other management-firewall rules.

accept-decisively

This option is functionally the same as accept.

drop Specifies that a matching packet should be silently dropped. The security software sends nothing back to the packet source, and it does not compare the packet to any other management-firewall rules.

reject

Specifies that a matching packet should be dropped. For TCP-based protocols, the security software sends a TCP reset (with the RST flag raised) back to the source. For other protocols, reject is equivalent to drop.

app-service

Associates the management-rule list with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers. The asm module has components for working with iApps.

description

Your description for the current rule.

destination

Matches against each packet's destination IP and/or destination port. The next options choose the matching criteria.

address-lists

Specifies a list of IP-address lists (see "security firewall address-list") to compare against the packet's destination address.

This list uses the same add, delete, none, and replace-all-with options described above for rules, as well as a default option.

addresses

Specifies a list of IP addresses and/or subnets to compare against the packet's destination address.

The format for an IPv4 address is a.b.c.d[/prefix]. The general format for an IPv6 address is a:b:c:d:e:f:g:h[/prefix]; you can shorten this by eliminating leading zeros from each field (for example, you can shorten "2001:0db7:3f4a:09dd:ca90:ff00:0042:8329" to "2001:db7:3f4a:9dd:ca90:ff00:42:8329"), and/or by removing the longest contiguous field of zeros (for example, you can shorten "2001:0:0:0:c34a:0:23ff:678" to "2001::c34a:0:23ff:678"). TMSH accepts any valid text representation of IPv6 addresses, as defined in RFC 2373 (see).

To edit this list, use the same add, delete, modify, none, and replace-all-with options described above for rules.

port-lists

Specifies a collection of port lists (see "security firewall port-list") to compare against the packet's destination port. If you use this option to specify a port list, a packet only matches if its destination port matches a port on these lists.

This list uses the same add, delete, none, and replace-all-with options described above for rules, as well as a default option.

ports

Specifies a list of ports and port ranges to compare against the packet's destination port.

To edit this list, use the same add, delete, modify, none, and replace-all-with options described above for rules.

icmp Specifies a list of ICMP types and codes to compare against the packet. You must set the ip-protocol option to "icmp" for this option to function. If you use this option, the current rule only matches ICMP packets that have the ICMP properties you specify here. You can add, delete, or modify (that is, change the description of) any entry in the list, or replace-all-with a new set of entries that you specify between curly braces ({}).

Use the standard integer identifiers to specify an ICMP type. For example: 3 is destination unreachable and 3:1 is destination unreachable with a code of host unreachable. The official list of ICMP types and codes is here: .

ip-protocol

Specifies the IP protocol to compare against the packet. This could be a layer-3 protocol (such as ipv4 or ipv6), or a higher-level protocol like ospf, rdp, or icmp. If you specify this option, a packet only matches if it uses the chosen protocol. Press the key for a full list of valid protocols.

log Specifies whether the security software should write a log entry for all packets that match this rule. You must also enable network filter logging in the "security log profile" component for this option to have any effect. Note that the security software always increments the statistics counter when a packet matches a rule, no matter how you set this option.

place-after [first | last | rule-name]

Specifies that a new rule should be placed after the first rule, the last rule, or the rule-name you

specify. If you are adding individual rules (as opposed to specifying replace-all-with), then you must use place-before or place-after to specify the rule's position in the list.

`place-before [first | last | rule-name]`

Specifies that a new rule should be placed before the first rule, the last rule, or the rule-name you specify. If you are adding individual rules (as opposed to specifying replace-all-with), then you must use place-before or place-after to specify the rule's position in the list.

`rule-list`

Specifies a full rule list instead of a customized rule that you might define with the other options. See "security firewall rule-list". If you use this option, then only the schedule and status options are valid; the tmsh software rejects any other options that you attempt to use with rule-list.

`schedule`

Specifies a schedule for the rule. See "security firewall schedule". If you omit this option, the rule or rule list is enabled all the time.

If the rule refers to a rule-list, the rule-list is enabled according to the schedule. When the rule list is enabled, the security software then honors the schedules defined within the rule-list.

`source`

Matches against each packet's source IP, source port, and/or source VLAN. The next options choose the matching criteria.

`address-lists`

Specifies a list of address lists (see "security firewall address-list") to compare against the packet's source address.

This list uses the same add, delete, none, and replace-all-with options described above for rules, as well as a default option.

`addresses`

Specifies a list of IP addresses and networks to compare against the packet's source address.

The format for an IPv4 address is a.b.c.d. The general format for an IPv6 address is a:b:c:d:e:f:g:h.

To edit this list, use the same add, delete, modify, none, and replace-all-with options described above for rules.

`port-lists`

Specifies a collection of port lists (see "security firewall port-list") to compare against the packet's source port. If you use this option to specify a port list, a packet only matches if its source port matches a port on these lists.

This list uses the same add, delete, none, and replace-all-with options described above for rules, as well as a default option.

`ports`

Specifies a list of ports and port ranges to compare against the packet's source port.

To edit this list, use the same add, delete, modify, none, and replace-all-with options described above for rules.

`vlan`

Specifies a list of VLANs, VLAN groups, and tunnels to compare against the packet.

This list uses the same add, delete, none, and replace-all-with options described above for rules, as well as a default option.

`status`

Specifies whether the rule is enabled, disabled or scheduled. A rule that is enabled is always checked. A rule that is disabled is never checked. A rule that is scheduled is checked according to the corresponding schedule configuration. A rule that is scheduled must have an associated schedule configuration.

`uid` Specifies how this rule UUID is assigned: assign an explicit uid based on RFC-4122, empty UUID (none value), or an auto-generated uid by system (auto-generated value) based on system wide mode:[uid-default-autogenerate mode] when creating a rule.

SEE ALSO

`cm config-sync`, `cm device-group`, `edit`, `list`, `modify`, `security firewall address-list`, `security firewall port-list`, `security firewall rule-list`, `security log profile`, `security firewall schedule`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015-2016. All rights reserved.

security firewall matching-rule

NAME

matching-rule - Shows the best match firewall rule amongst all the admin configured Network Firewall rules in different contexts (global, route-domain, VIP/SelfIP) given source/destination IP address and port, protocol and user configured vlan name. You can only use the show command with this component.

MODULE

security firewall

SYNTAX

```
show matching-rule
  dest-addr [IP address]
  source-addr [IP address]
  dest-port [TCP/UDP port]
  source-port [TCP/UDP port]
  protocol [protocol]
  vlan [vlan name]
```

DESCRIPTION

With user provided VLAN, source/destination IP addresses, TCP/UDP ports and protocol, the command will try to match these parameters against user configured ACL rules in global, route domain, VIP/SelfIP context, and return the best match rules. Both IPv4 and IPv6 addresses and all possible protocols are supported. This command can be used as a diagnostic tool to trouble-shoot BigIP firewall configuration problem. It provides a faster way to identify which ACL rule will have impact to the specified packet stream.

EXAMPLES

```
# show security firewall matching-rule dest-addr 1.1.1.1 dest-port 140 source-addr 2.2.2.2 source-port 141
protocol 10 vlan /Common/internal
```

Firewall Matching Rule:

```
-----
Context Type Context Name Policy Name Rule Name Action
-----
Global    globalrule Accept
```

Total records returned: 1

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-09 security firewall matching-rule(1)

security firewall on-demand-compilation

NAME

on-demand-compilation - Configures the compilation mode of firewall rules.

MODULE

security firewall

SYNTAX

```
MODIFY
modify on-demand-compilation mode [disabled | enabled]
```

DISPLAY

```
list on-demand-compilation mode
show running-config on-demand-compilation mode
```

TRIGGER ON-DEMAND COMPILATION

```
run on-demand-compilation
```

DESCRIPTION

You can use the on-demand-compilation component to change the behavior of firewall rule compilation. By default the system will automatically compile the changes at each configuration change. You can change the behavior to manually trigger the compilation by setting the mode to enabled.

EXAMPLES

```
modify security firewall on-demand-compilation mode enabled
```

run security firewall on-demand-compilation

SEE ALSO

security firewall on-demand-rule-deploy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-09-16 security firewall on-demand-compilation(1)

security firewall on-demand-rule-deploy

NAME

on-demand-rule-deploy - Configures the rule deploy mode of firewall rules.

MODULE

security firewall

SYNTAX

MODIFY

modify on-demand-rule-deploy mode [disabled | enabled]

DISPLAY

list on-demand-rule-deploy mode

show running-config on-demand-rule-deploy mode

TRIGGER ON-DEMAND RULE DEPLOY

run on-demand-rule-deploy

DESCRIPTION

You can use the on-demand-rule-deploy component to change the behavior of firewall rule deployment. By default the system will automatically deploy the changes once the compilation is successful. You can change the behavior to manually trigger the rule deployment by setting the mode to enabled.

EXAMPLES

modify security firewall on-demand-rule-deploy mode enabled

run security firewall on-demand-rule-deploy

SEE ALSO

security firewall on-demand-compilation

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2017-08-18 security firewall on-demand-rule-deploy(1)

security firewall policy

NAME

policy - Configures firewall policy.

MODULE

security firewall

SYNTAX

Modify the policy component within the security firewall module using the syntax shown in the following sections.

CREATE/MODIFY

create policy [name]

options:

copy-from [string]

```

modify policy [name]
options:
  description [string]
  rules [add | delete | modify | replace-all-with] {
    [ [name] ] {
      options:
action [accept | accept-decisively | drop | reject]
description [string]
destination {
  address-lists [add | default | delete | replace-all-with] {
    [address list names...]
  }
  address-lists none
  addresses [add | default | delete | replace-all-with] {
    [ [ip address] | [ip address/prefixlen] ]
  }
  addresses none
  fqdns [add | delete | replace-all-with] {
    [ fully qualified domain names]
  }
  fqdns none
  geo [add | default | delete | replace-all-with] {
    [ [country_code [state state_name] ] ]
  }
  geo none
  ipi-category [add | default | delete | replace-all-with] {
    [ IP-Intelligence category names... ]
  }
  ipi-category none
  port-lists [add | default | delete | replace-all-with] {
    [port list names...]
  }
  port-lists none
  ports [add | default | delete | none | replace-all-with] {
    [ [port] | [port1-port2] ]
  }
  ports none
  zones [add | delete | replace-all-with] {
    [ zone names]
  }
  zones none
}
icmp [add | delete | modify | replace-all-with] {
  [ [icmp_type] | icmp_type:icmp_code ] {
    description [string]
  }
}
icmp none
ip-protocol [protocol name]
irule [irule name]
irule-sample-rate [integer]
log [no | yes]
place-after [first | last | [rule name]]
place-before [first | last | [rule name]]
rule-list [rule list name]
schedule [schedule name]
uuid [ | none | auto-generate]
source {
  address-lists [add | default | delete | replace-all-with] {
    [address list names...]
  }
  address-lists none
  addresses [add | default | delete | replace-all-with] {
    [ [ip address] | [ip_address/prefixlen] ]
  }
  addresses none
  fqdns [add | delete | replace-all-with] {
    [ fully qualified domain names]
  }
  fqdns none
  geo [add | default | delete | replace-all-with] {
    [ [country_code [state state_name] ] ]
  }
  geo none
  identity {
    user-groups [add | delete | modify | none | replace-all-with] {
[user group names...]
  }
    user-lists [add | delete | modify | none | replace-all-with] {
[user list names...]
  }
    users [add | delete | modify | none | replace-all-with] {
[user names...]
  }
  }
  ipi-category [add | default | delete | replace-all-with] {
    [ IP-Intelligence category names... ]
  }
}

```

```

}
ipi-category none
port-lists [add | default | delete | replace-all-with] {
  [port list names...]
}
port-lists none
ports [add | default | delete | replace-all-with] {
  [ [port] | [port1-port2] ]
}
ports none
vlans [add | default | delete | replace-all-with] {
  [vlan names...]
}
vlans none
zones [add | delete | replace-all-with] {
  [ zone names]
}
zones none
}
status [disabled | enabled | scheduled]
service-policy [service policy name]
virtual-server [virtual server name]
ips-profile [IPS profile name]
classification-policy [classification policy name]
}
}
rules none

```

```

edit policy
options:
  all-properties
  non-default-properties

```

```

DISPLAY
list policy
show running-config policy
options:
  all-properties
  non-default-properties
  one-line

```

DESCRIPTION

You can use the policy component to configure a shareable and reusable set of network firewall rules which can be associated as enforced or staged with a number of configuration objects of the following types: net self, ltm virtual, security firewall global-rules, net route-domain.

EXAMPLES

```

modify policy rules add {
reject-internal-net {
  place-before first
  action reject
  source {
    addresses replace-all-with { 172.27.0.0/16 }
  }
}
}

```

Creates a rule entry at the beginning of the list that rejects traffic from the 172.27.0.0 network.

```
modify policy rules delete reject-internal-net
```

Removes the rule reject-internal-net from the list of rules.

```
create security firewall policy p1 rules add { r1 { source { geo add { US } } action reject place-after first } }

```

Creates a policy with a single rule that rejects all packets from the US.

```
create security firewall policy xyz rules add { r1 { destination { fqdns add { f5.com } } action accept place-after first } }

```

Creates a policy named 'xyz' with a single rule (named 'r1') that accepts all packets with destination IP address in domain 'f5.com'.

```
list policy
```

Displays the current list of policy rules.

```
create policy "New Policy" copy-from "/Common/Existing Policy"
```

Creates a new policy New Policy by copying existing policy /Common/Existing Policy.

OPTIONS

description
User defined description.

copy-from
(CREATE)Specifies the name of an existing policy from which to copy all configuration options.

rules
Adds, deletes, or replaces a firewall rule.

action

Specifies the action that the system takes when a rule is matched.

accept

Specifies that the current packet should be accepted.

accept-decisively

Specifies that the current packet should be accepted and that packet will not be compared to any other firewall rules in any other context.

drop Specifies that the current packet should be silently dropped. Nothing is sent back to the packet source. The packet is not compared to any other firewall rules.

reject

Specifies that the current packet should be dropped. For TCP based protocols a TCP reset is sent to the source. For other protocols reject is equivalent to drop.

description

User defined description.

destination

address-lists

Specifies a list of address lists (see security firewall address-list) against which the packet will be compared.

addresses

Specifies a list of addresses and networks against which the packet will be compared.

fqdns

Specifies a list of fully qualified domain names to compare against packet's destination IP address domain.

geo Specifies a list of Geo Locations that the packet will be compared against.

ipi-category

Specifies a list of IP-Intelligence category names that the packet will be compared against.

port-lists

Specifies a list of port lists (see security firewall port-list) against which the packet will be compared.

ports

Specifies a list of ports and port ranges against which the packet will be compared.

zones

Specifies a list of zones, (see security firewall zone) against which the packet will be compared.

icmp Specifies a list of ICMP types and codes against which the packet will be compared. The standard integer identifiers are used to specify an ICMP type Example: 3 is destination unreachable and 3:1 is destination unreachable with a code of host unreachable. The list of ICMP types and codes can be found here <http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xml>.

ip-protocol

Specifies the IP protocol against which the packet will be compared.

irule

Specifies the name of the iRule that will be triggered when a packet matches this firewall rule. The firewall rule match raises a FLOW_INIT iRule event.

irule-sample-rate

Specifies the rate at which an iRule specified by irule option will be triggered when a packet matches this firewall rule. The rate is an integer value in the range 0-65535 and specifies how many packets must match this firewall rule before the iRule is triggered. The default value is 1 and causes the iRule to be triggered for every packet that matches this firewall rule. A value of 0 disables iRule triggering.

log Specifies whether the packet will be logged if it matches the rule. Logging must also be enabled in the corresponding logging configuration. (e.g. security log profile global-network when policy assigned to global-rules). Note that the statistics counter is always incremented when a packet matches a rule.

place-after

Specifies that a new rule should be placed after another rule, first or last. If individual rules are being added (as opposed to specifying replace-all-with) then place-before or place-after must be specified.

place-before

Specifies that a new rule should be placed before another rule, first or last. If individual rules are being added (as opposed to specifying replace-all-with) then place-before or place-after must be specified.

rule-list

Specifies a list of rules to evaluate. See security firewall rule-list. If a rule-list is specified then only the schedule and status properties effect the rule.

schedule

Specifies a schedule for the rule. See security firewall schedule. If the rule refers to a rule-

list the rule-list will be enabled according to the schedule. When the rule list is enabled, the schedules defined within the rule-list will be honored.

source
address-lists

Specifies a list of address lists (see security firewall address-list) against which the packet will be compared.

addresses

Specifies a list of addresses and networks against which the packet will be compared.

fqdns

Specifies a list of fully qualified domain names to compare against packet's source IP address domain.

geo Specifies a list of Geo Locations against which the packet will be compared.

ipi-category

Specifies a list of IP-Intelligence category names that the packet will be compared against.

port-lists

Specifies a list of port lists (see security firewall port-list) against which the packet will be compared.

ports

Specifies a list of ports and port ranges against which the packet will be compared.

vlan

Specifies a list of vlans, vlan groups and tunnels against which the packet will be compared.

zones

Specifies a list of zones, (see security firewall zone) against which the packet will be compared.

status

Specifies whether the rule is enabled, disabled or scheduled. A rule that is enabled is always checked. A rule that is disabled is never checked. A rule that is scheduled is checked according to the corresponding schedule configuration. A rule that is scheduled must have an associated schedule configuration.

service-policy

Specifies the service policy configuration to use. (see "net service-policy"). The service policy can be used to set specific policy based configurations like flow timers, which applies to the flows that matches the rule.

uuid Specifies how this rule UUID is assigned: assign an explicit uuid based on RFC-4122, empty UUID (none value), or an auto-generated uuid by system (auto-generated value) based on system wide mode:[uuid-default-autogenerate mode] when creating a rule.

virtual-server

Specifies the virtual server name that will be used for further traffic processing. Option is valid only for global and/or route domain contexts.

SEE ALSO

create, edit, list, modify, security firewall address-list, security firewall port-list, security firewall rule-list, security log profile, security firewall schedule, net service-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

BIG-IP 2018-09-17 security firewall policy(1)

security firewall port-list

NAME

port-list - Configures a port-list for use by firewall rules. A firewall rule can match a packet's source port or destination port against one of the ports in a port list, and can take some action (such as ACCEPT or DROP) for a matching packet.

MODULE

security firewall

SYNTAX

CREATE/MODIFY
create port-list [name]

modify port-list [[name] | all]

options:

app-service [name]

description [string]

ports [add | delete | modify | replace-all-with] {
[[port] | [port] - [port]]
}

edit port-list [[name] | all]

options:

all-properties

non-default-properties

DISPLAY

list port-list [[name] | all | [property]]

show running-config port-list [[name] | all | [property]]

DELETE

delete port-list [[name] | all]

DESCRIPTION

You can use the port-list component to define reusable lists of ports for various firewall rules. The network software compares a packet's source port and/or destination port against ports in this list. You can assign a port list to the firewall rules in net self, net route-domain, security firewall global-rules, security firewall rule-list, sys management-ip, and Itm virtual firewall rules.

EXAMPLES

```
create port-list p-list1 ports add { 80 }
```

Creates a new port list with one entry.

```
list port-list
```

```
security firewall port-list _sys_self_allow_tcp_defaults {
```

```
  ports {  
    domain { }  
    f5-iquery { }  
    https { }  
    snmp { }  
    ssh { }  
  }  
}
```

```
}
```

```
security firewall port-list _sys_self_allow_udp_defaults {
```

```
  ports {  
    520 { }  
    cap { }  
    domain { }  
    f5-iquery { }  
    snmp { }  
  }  
}
```

```
}
```

```
security firewall port-list p-list1 {
```

```
  ports {  
    http { }  
  }  
}
```

```
}
```

Shows all the port lists, including the one created in the previous example.

OPTIONS

app-service

Associates this port list with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers. The asm module has components for working with iApps.

description

Your description for the port list.

ports

Specifies a list of ports to compare against a packet's source or destination port. Use one of the keywords below and then specify the port(s) to add or delete. Specify ranges of ports with a dash between the two ends of the range (for example, 80-88).

add Creates a new port list, which you specify next with port numbers in curly braces ({}).

delete

Deletes the port(s) that you specify next, in curly braces ({}).

modify

Is not supported for this component.

replace-all-with

Replaces the current set of ports with the port(s) that you specify next, in curly braces ({}).

SEE ALSO

edit, list, modify, net self, net route-domain, security firewall address-list, security firewall rule-list, security firewall global-rules, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 security firewall port-list(1)

security firewall port-misuse-policy

NAME

port-misuse-policy - Configures the port misuse policies.

MODULE

security firewall

SYNTAX

Configure the port misuse policy component within the security firewall module using the syntax shown in the following sections.

CREATE/MODIFY

```
create port-misuse-policy [name]
```

```
modify port-misuse-policy [name]
```

options:

```
app-service [[string] | none]
```

```
description [string]
```

```
drop-on-l7-mismatch [no | yes]
```

```
log-on-l7-mismatch [no | yes]
```

```
rules [add | delete | modify | replace-all-with] {
```

```
  [rule name] }
```

options:

```
description [string]
```

```
drop-on-l7-mismatch [no | yes | use-policy-setting]
```

```
ip-protocol [sctp | tcp | udp]
```

```
l7-protocol [protocol name]
```

```
log-on-l7-mismatch [no | yes | use-policy-setting]
```

```
port [port]
```

```
}
```

```
}
```

```
rules none
```

```
edit port-misuse-policy [[name] | all]
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list port-misuse-policy
```

```
show running-config port-misuse-policy
```

options:

```
all-properties
```

```
non-default-properties
```

```
one-line
```

DELETE

```
delete port-misuse-policy [[name] | all]
```

DESCRIPTION

You can use the port-misuse-policy component to configure a shareable and reusable set of network port misuse policies which can be associated with a service policy objects. A port misuse policy has one or more rules that match connections by IP transport layer (L4) protocol and port number. Each rule must have a unique L4 protocol and port combination within the policy. When connection matches a policy rule (i.e. L4 protocol/port pair) the first data packet of the connection is tested to conform to application (L7) protocol specified in the rule. If data conforms to the L7 protocol (or test is inconclusive) the policy stops. The connection is allowed to proceed normally and data is processed as if no policy is in use. If data definitely does not conform to the specified L7 protocol the connection is treated according to configuration of the matched rule or the policy if rule uses policy defaults. In this case the rule or policy can drop the connection or allow it to proceed, and can also log an event about L7 protocol mismatch.

Port misuse policy (via service policy) can be associated with objects of the following types: ltm virtual, net route-domain, global. Several port misuse policies can be associated with objects of each type. In addition to service policy specified in the object itself, service policies could be associated with ACL rules of the security firewall policy, if it is associated with the object. When more than one policy is associated with the object the most specific port misuse rule is used. For example, if connection matches an ACL rule which has a service policy with port misuse policy also having a rule matching the connection, that port misuse rule is applied. Otherwise the port misuse rule associated via virtual's service policy is applied, if such rule exists and matches the connection. See also net service-policy.

Port misuse policies could be specified for both virtual server and route domain objects associated with the connection. In this case all policies are applied. If a policy has a matching rule that drops the connection, and connection fails to pass L7 protocol test, the connection is terminated and remaining policies are not applied.

EXAMPLES

```
create security firewall port-misuse-policy web-ports-policy drop-on-l7-mismatch no log-on-l7-mismatch yes
rules add { p80 { ip-protocol tcp port 80 l7-protocol http drop-on-l7-mismatch yes } p8080 { ip-protocol tcp
port 8080 l7-protocol http } }
```

```
list security firewall port-misuse-policy web-ports-policy
security firewall port-misuse-policy web-ports-policy {
  drop-on-l7-mismatch no
  log-on-l7-mismatch yes
  rules {
    p80 {
      drop-on-l7-mismatch yes
      l7-protocol http
      port http
    }
    p8080 {
      l7-protocol http
      port webcache
    }
  }
}
```

Creates port misuse policy with rules for tcp ports 80 and 8080 that test if first data packet looks like HTTP. The rule p80 tests all connections that have destination port TCP 80 and drops them if the first data packet does not look like HTTP. The rule p8080 tests all connections that have destination port TCP 8080 and logs an event if the first data packet does not look like HTTP (because of policy defaults).

```
modify security firewall port-misuse-policy web-ports-policy { rules add { p8888 { port 8888
drop-on-l7-mismatch yes } } }
```

Adds a new rule p8888 to port misuse policy web-ports-policy that tests all connections to TCP port 8888 and drops them and logs an event when the first data packet does not look like HTTP.

```
list security firewall port-misuse-policy
```

Displays the current port misuse policy configuration list.

OPTIONS

description
User defined description.

drop-on-l7-mismatch
Indicates if the connection should be dropped when there is a matching rule in the policy that has drop-on-l7-mismatch set to use-policy-setting and connection that matches that rule fails L7 protocol test. The default is yes.

log-on-l7-mismatch
Indicates if a port misuse event should be logged when there is a matching rule in the policy that has log-on-l7-mismatch set to use-policy-setting and connection that matches that rule fails L7 protocol test. The default is no.

rules
Adds, deletes, or replaces a named port misuse policy rule.

description
User defined description.

drop-on-l7-mismatch
Indicates if the connection should be dropped when it matches this rule but fails L7 protocol test. Allowed values are yes, no, and use-policy-setting. The default is use-policy-setting.

ip-protocol
Specifies the transport layer (L4) IP protocol for matching the connection. The valid protocols are sctp, tcp, and udp. A port and L4 protocol combination must be unique for the policy. The default is tcp.

l7-protocol
Specifies the application layer (L7) protocol for the rule. When the connection matches the rule the first data packet is tested to conform to this protocol. If the test is negative the rule can drop the connection and/or log a port misuse event depending on other options. If the test is positive or inconclusive (not enough data) the connection is handled as if there was no port misuse policy associated with the given object (virtual server or route domain), and policies at other objects are applied. Press the `?` key for a full list of valid protocols. The default protocol is http.

log-on-l7-mismatch
Indicates if a port misuse event should be logged when the connection matches this rule but fails L7 protocol test. Allowed values are yes, no, and use-policy-setting. The default is use-policy-setting.

port Specifies the destination port number for matching the connection. The valid values are 1-65535. A port and L4 protocol combination must be unique for the policy.

SEE ALSO

create, edit, list, modify, security firewall rule-list, security firewall policy, net service-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2015. All rights reserved.

BIG-IP 2015-07-14 security firewall port-misuse-policy(1)

security firewall rule-list

NAME

rule-list - Configures a rule-list of network firewall rules. You can reuse a rule list in multiple firewalls, such as the firewalls for self IPs, routing domains, and the global firewall.

MODULE

security firewall

SYNTAX

```
CREATE/MODIFY
create rule-list [name]
modify rule-list [[name] | all]
options:
description [string]
rules [add | delete | modify | replace-all-with] {
  [ name ] {
    options:
action [accept | accept-decisively | drop | reject]
app-service [name]
description [string]
source {
address-lists [add | default | delete | replace-all-with] {
  [address list names...]
}
address-lists none
addresses [add | delete | modify | replace-all-with] {
  [ ip address ] | [ip_address/prefixlen] ]
}
addresses none
fqdns [add | delete | replace-all-with] {
  [ fully qualified domain names]
}
fqdns none
geo [add | default | delete | replace-all-with] {
  [ [country_code [state state_name] ] ]
}
geo none
ipi-category [add | default | delete | replace-all-with] {
  [ IP-Intelligence category names... ]
}
ipi-category none
port-lists [add | default | delete | replace-all-with] {
  [port list names...]
}
port-lists none
ports [add | default | delete | replace-all-with] {
  [ [port] | [port1-port2] ]
}
ports none
vlans [add | default | delete | replace-all-with] {
  [vlan names...]
}
vlans none
}
destination {
address-lists [add | default | delete | replace-all-with] {
  [address list names...]
}
address-lists none
addresses [add | delete | modify | replace-all-with] {
  [ ip address ] | [ip address/prefixlen] ]
}
addresses none
fqdns [add | delete | replace-all-with] {
  [ fully qualified domain names]
}
}
```

```

fqdns none
geo [add | default | delete | replace-all-with] {
  [ [country_code [state state_name] ] ]
}
geo none
ipi-category [add | default | delete | replace-all-with] {
  [ IP-Intelligence category names... ]
}
ipi-category none
port-lists [add | default | delete | replace-all-with] {
  [port list names...]
}
port-lists none
ports [add | delete | modify | replace-all-with] {
  [ [port] | [port1-port2] ]
}
ports none
}
icmp [add | delete | modify | replace-all-with] {
  [ [icmp_type] | icmp_type:icmp_code ] {
    description [string]
  }
}
icmp none
ip-protocol [protocol name]
irule [irule name]
irule-sample-rate [integer]
log [no | yes]
place-after [first | last | [rule name]]
place-before [first | last | [rule name]]
rule-list [rule list name]
schedule [schedule name]
status [disabled | enabled | scheduled]
service-policy [service policy name]
uuid [ | none | auto-generate]
virtual-server [virtual server name]
ips-profile [IPS profile name]
classification-policy [classification policy name]
}
}
rules none

edit rule-list [[name] | all]
options:
  all-properties
  non-default-properties

DISPLAY
list rule-list [[name] | all | [property]]
show running-config rule-list [[name] | all | [property]]

```

DESCRIPTION

You can use the rule-list component to configure network firewall rules to be applied to multiple firewalls. The network software compares IP packets to the criteria specified in these rules. If a packet matches the criteria then the system takes the action specified by the rule. If a packet does not match any rule in the list, the software accepts the packet or passes it to the next rule or rule-list (for example, the system compares the packet to net self-ip rules if the packet is destined for a network associated with a self-ip that has firewall rules defined).

MATCHING AN IP PACKET

You can use this TMSH component to match against any or all of the following properties of an IP packet:

```

source address
source fqdn
source geo
source port
the packet's source VLAN
destination address
destination fqdn
destination geo
destination port
the higher-level protocol in the packet's payload

```

If you match against more than one of these items, a packet must pass all of your tests to successfully match. For example, if you match against a source subnet and several destination ports, a packet must originate from the given subnet and must also have one of the specified destination ports.

RULE ORDER

The network software evaluates firewall rules in the order that you specify. You can use the list management-ip-rules command to see the current rule order. As you add or modify rules in this component, you can use the place-before rule-name or place-after rule-name option to choose the rule's place in the sequence.

Rule order can determine whether or not a packet is dropped. Consider the following rules:

```

rule_a, matches source addresses against 172.16.0.0 and ACCEPTS all packets that match.
rule_d, matches source addresses against 172.16.39.0 and DROPS all packets that match.

```

Also consider a packet from a host at 172.16.39.55. If rule_a appears before rule_d in the rule list, the packet's source address matches rule_a first and the software accepts it. The software never reaches rule_d for comparison. If rule_d appears first instead, the packet's source address now matches rule_d; in this case, the software drops the packet.

EXAMPLES

```
create rule-list block_bad_mgmt description "ports to be blocked on our management interfaces" rules replace-all-with { reject_telnet { ip-protocol tcp destination { ports add { telnet } } action reject } }
```

Creates a new rule list called block_bad_mgmt. It matches and rejects any TCP packet whose destination port is telnet. The description indicates that the rule is intended for the management-IP firewall.

```
modify rule-list block_bad_mgmt rules add { reject_http { ip-protocol tcp destination { ports add { http } } action reject place-after last } }
```

Modifies the above rule list by blocking HTTP traffic, too.

```
list rule-list block_bad_mgmt
security firewall rule-list block_bad_mgmt {
  description "ports to be blocked on our management interfaces"
  rules {
    reject_telnet {
      action reject
      destination {
        ports {
          telnet { }
        }
      }
    }
    ip-protocol tcp
  }
  reject_http {
    action reject
    destination {
      ports {
        http { }
      }
    }
  }
  ip-protocol tcp
}
}
```

Shows the above rule list, with both rules.

```
modify rule-list rules add { reject-internal-net { place-before first action reject source { addresses replace-all-with { 172.27.0.0/16 } } } }
```

Creates a rule entry at the beginning of the list that rejects traffic from the 172.27.0.0 network.

```
create security firewall rule-list r1 description "Geo Locations to be blocked" rules add { r1 { source { geo add { US } } place-after first action drop } }
```

Creates a new rule list "r1", which matches and rejects any packet with a US source. The description explains the purpose of the rule list.

```
modify security firewall rule-list r1 rules add { r2 { source { geo add { CA } } place-before last action drop } }
```

```
security firewall rule-list r1 {
  description "Geo Locations to be blocked"
  rules {
    r2 {
      action drop
      source {
        geo {
          CA {
            state none
          }
        }
      }
    }
    r1 {
      action drop
      source {
        geo {
          US {
            state none
          }
        }
      }
    }
  }
}
```

Shows the above rule list, with both rules.

```
create security firewall rule-list r1 description "domains to be blocked" rules add { r1 { destination {
```

```
fqdns add { xyz.com} } place-after first action drop } }
```

Creates a new rule list "r1", which matches and rejects any packet with destination IP addresses in domain 'xyz.com'. The description explains the purpose of the rule list.

```
modify security firewall rule-list r1 rules modify { r1 { destination { fqdns add { abc.com } } } }
```

Modifies the above rule list by blocking destination IP addresses in domain 'abc.com' too.

```
list rule-list r1
security firewall rule-list r1 {
description "domains to be blocked"
rules {
r1 {
action drop
destination {
fqdns {
abc.com { }
xyz.com { }
}
}
}
}
}
```

Shows the above rule list, with the single rule r1.

OPTIONS

app-service

Associates the rule list with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers. The asm module has components for working with iApps.

description

Your description for this list of firewall rules.

rules

Adds, deletes, or replaces a firewall rule.

add Creates a new rule, which you specify next with a unique string in curly braces ({}). Use the place-before or place-after option inside the curly braces to determine the order of the rule. If this is the first rule, use the replace-all-with option instead of add.

delete

Deletes the rule that you specify next, in curly braces ({}).

modify

Modifies the existing rule that you specify next, in curly braces ({}). After the rule name, enter the new configuration settings for the rule inside a nested set of curly braces.

none Empties the list of rules. An empty rule list implicitly accepts all packets. The security software skips this context and assesses packets against the next layer of firewall rules, if there is one (such as those defined for net self-ip, net route-domain or ltm virtual)

replace-all-with

Replaces the current list of rules with the rule(s) that you specify next, in curly braces ({}). Use this option for the first rule in the list.

Enter the name of a rule to be added or modified, then enter an open curly brace ({}), one or more of the following options, and a closed curly brace ({}).

action

Specifies the action that the system takes when a rule is matched.

accept

Specifies that a matching packet should be accepted. The security software stops comparing a matching packet to any other rules in the list. The software continues comparing the packet to rules in the next appropriate context (such as net self-ip, net route-domain or ltm virtual).

accept-decisively

Specifies that a matching packet should be accepted and should not be compared to any other firewall rules in any other context.

drop Specifies that a matching packet should be silently dropped. The security software sends nothing back to the packet source. The security software does not compare the packet to any other firewall rules in any other context.

reject

Specifies that a matching packet should be dropped. For TCP-based protocols, the security software sends a TCP reset (with the RST flag raised) back to the source. For other protocols, reject is equivalent to drop.

description

Your description for the current rule.

destination

Matches against each packet's destination IP and/or destination port. The next options choose the

matching criteria.

address-lists

Specifies a list of IP-address lists (see "security firewall address-list") to compare against the packet's destination address.

This list uses the same add, delete, none, and replace-all-with commands described above for rules, as well as a default command.

addresses

Specifies a list of IP addresses and/or subnets to compare against the packet's destination address.

The format for an IPv4 address is a.b.c.d[/prefix]. The general format for an IPv6 address is a:b:c:d:e:f:g:h[/prefix]; you can shorten this by eliminating leading zeros from each field (for example, you can shorten "2001:0db7:3f4a:09dd:0a90:ff00:0042:8329" to "2001:db7:3f4a:9dd:a90:ff00:42:8329"), and/or by removing the longest contiguous field of zeros (for example, you can shorten "2001:0:0:0:c34a:0:0:678" to "2001::c34a:0:0:678"). TMSH accepts any valid text representation of IPv6 addresses, as defined in RFC 2373 (see).

To edit this list, use the same add, delete, modify, none, and replace-all-with commands described above for rules.

fqdns

Specifies a list of fully qualified domain names to compare against packet's destination IP address domain.

To edit this list, use the same add, delete, none, and replace-all-with commands described above for rules.

geo Specifies a list of Geo Locations to compare a packet's source or destination Geo Location.

The format for a Geo Location is a 2 character string for the country code and a string for the state.

To edit this list, use the same add, delete, modify, none, and replace-all-with options described above for rules.

ipi-category

Specifies a list of IP-Intelligence category names that the packet will be compared against.

port-lists

Specifies a collection of port lists (see "security firewall port-list") to compare against the packet's destination port. If you use this option to specify a port list, a packet only matches if its destination port matches a port on these lists.

If you combine address lists and port lists in the same rule, a packet must have a matching port and a matching address to fully match the rule.

This list uses the same add, delete, none, and replace-all-with commands described above for rules, as well as a default command.

ports

Specifies a list of ports and port ranges to compare against the packet's destination port.

To edit this list, use the same add, delete, modify, none, and replace-all-with commands described above for rules.

icmp Specifies a list of ICMP types and codes to compare against the packet. You must set the ip-protocol option to "icmp" for this option to function. If you use this option, the current rule only matches ICMP packets that have the ICMP properties you specify here. You can add, delete, or modify (that is, change the description of) any entry in the list, or replace-all-with a new set of entries that you specify between curly braces ({}).

Use the standard integer identifiers to specify an ICMP type. For example: 3 is destination unreachable and 3:1 is destination unreachable with a code of host unreachable. The official list of ICMP types and codes is here: .

ip-protocol

Specifies the IP protocol to compare against the packet. This could be a layer-3 protocol (such as ipv4 or ipv6), or a higher-level protocol like ospf or rdp. If you specify this option, a packet only matches if it uses the chosen protocol. Press the key for a full list of valid protocols.

irule

Specifies the name of the iRule that will be triggered when a packet matches this firewall rule. The firewall rule match raises a FLOW_INIT iRule event.

irule-sample-rate

Specifies the rate at which an iRule specified by irule option will be triggered when a packet matches this firewall rule. The rate is an integer value in the range 0-65535 and specifies how many packets must match this firewall rule before the iRule is triggered. The default value is 1 and causes the iRule to be triggered for every packet that matches this firewall rule. A value of 0 disables iRule triggering.

log Specifies whether the security software should write a log entry for all packets that match this rule. You must also enable network filter logging in the "security log profile" component for this

option to have any effect. Note that the security software always increments the statistics counter when a packet matches a rule, no matter how you set this option.

place-after [first | last | rule-name]

Specifies that a new rule should be placed after the first rule, the last rule, or the rule-name you specify. If you are adding individual rules (as opposed to specifying replace-all-with), then you must use place-before or place-after to specify the rule's position in the list.

place-before [first | last | rule-name]

Specifies that a new rule should be placed before the first rule, the last rule, or the rule-name you specify. If you are adding individual rules (as opposed to specifying replace-all-with), then you must use place-before or place-after to specify the rule's position in the list.

rule-list

Specifies a full rule list instead of a customized rule that you might define with the other options. If you use this option, then only the schedule and status options are valid; the tmsh software rejects any other options that you attempt to use with rule-list.

schedule

Specifies a schedule for the rule. See "security firewall schedule". If you omit this option, the rule or rule list is enabled all the time.

If the rule refers to a rule-list, the rule-list is enabled according to the schedule. When the rule list is enabled, the security software then honors any schedules defined within the rule-list.

source

Matches against each packet's source IP, source port, and/or source VLAN. The next options choose the matching criteria.

address-lists

Specifies a list of address lists (see "security firewall address-list") to compare against the packet's source address.

This list uses the same add, delete, none, and replace-all-with commands described above for rules, as well as a default command.

addresses

Specifies a list of IP addresses and networks to compare against the packet's source address.

The format for an IPv4 address is a.b.c.d. The format for an IPv6 address is a:b:c:d:e:f:g:h.

To edit this list, use the same add, delete, modify, none, and replace-all-with commands described above for rules.

fqdns

Specifies a list of fully qualified domain names to compare against packet's source IP address domain.

To edit this list, use the same add, delete, none, and replace-all-with commands described above for rules.

geo Specifies a list of Geo Locations to compare a packet's source or destination Geo Location.

The format for a Geo Location is a 2 alphabet string for the country code and a string for the state.

To edit this list, use the same add, delete, modify, none, and replace-all-with options described above for rules.

ipi-category

Specifies a list of IP-Intelligence category names that the packet will be compared against.

port-lists

Specifies a collection of port lists (see "security firewall port-list") to compare against the packet's source port. If you use this option to specify a port list, a packet only matches if its source port matches a port on these lists.

This list uses the same add, delete, none, and replace-all-with commands described above for rules, as well as a default command.

ports

Specifies a list of ports and port ranges to compare against the packet's source port.

To edit this list, use the same add, delete, modify, none, and replace-all-with commands described above for rules.

vlan

Specifies a list of VLANs, VLAN groups, and tunnels to compare against the packet.

This list uses the same add, delete, none, and replace-all-with commands described above for rules, as well as a default command.

status

Specifies whether the rule is enabled, disabled or scheduled. A rule that is enabled is always checked. A rule that is disabled is never checked. A rule that is scheduled is checked according to the corresponding schedule configuration. A rule that is scheduled must have an associated schedule configuration.

service-policy

Specifies the service policy configuration to use. (see "net service-policy"). The service policy can be used to set specific policy based configurations like flow timers, which applies to the flows that matches the rule.

uuid Specifies how this rule UUID is assigned: assign a explicit uuid based on RFC-4122, empty UUID (none value), or an auto-generated uuid by system (auto-generated value) based on system wide mode:[uuid-default-autogenerate mode] when creating a rule.

virtual-server

Specifies the virtual server name that will be used for further traffic processing. Option is valid only for global and/or route domain contexts.

ips-profile

Specifies IPS profile name used for signature matching and/or protocol compliance checks for flows matching the rule.

classification-policy

Specifies the Classification Policy name that will be enforced.

SEE ALSO

edit, list, modify, security firewall address-list, security firewall port-list, security firewall global-rules, security log profile, security firewall schedule, net service-policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

BIG-IP 2018-09-17 security firewall rule-list(1)

security firewall rule-stat

NAME

rule-stat - Displays statistics of firewall rules on the BIG-IP(r) system. You can only use the show command with this component.

MODULE

security firewall

SYNTAX

show rule-stat

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

DESCRIPTION

You can use the rule-stat component to display statistics of firewall rules.

EXAMPLES

show rule-stat

Displays firewall rule's statistics in the system default units.

show rule-stat raw

Displays raw firewall rule's statistics.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-03-21 security firewall rule-stat(1)

security firewall schedule

NAME

schedule - Create a schedule that you can apply to firewall rules.

MODULE

security firewall

SYNTAX

CREATE/MODIFY

create schedule [name]

modify schedule [[name] | all]

options:

app-service [name]

daily-hour-end [hour:minute]

daily-hour-start [hour:minute]

date-valid-end [MM/DD/YYYY]

date-valid-start [MM/DD/YYYY]

description [text]

days-of-week [[monday | tuesday | wednesday | thursday | friday | saturday | sunday] ...]

edit schedule [[name] | [glob] | [regex]] ...]

DISPLAY

list schedule [[name] | all | [property]]

DELETE

delete schedule [[name] | all]

show running-config schedule [[name] | all | [property]]

DESCRIPTION

You use the schedule component to specify when to apply a firewall rule. You can specify a start time and an end time, some days of the week, a date when the schedule first starts, and/or a date when the schedule ends forever.

To apply the schedule to a firewall rule or rule list, edit the firewall or rule-list component. These are the firewalls and rule lists where you can apply schedules:

"security firewall global-rules"

"security firewall management-ip-rules"

"net self"

"net route-domain"

"itm virtual"

"security firewall rule-list"

By default, all firewall rules are continuously active. By applying a schedule to a firewall rule, you reduce the time that the rule is running.

If you create a schedule without any scheduling specifications (such as daily-hour-start), the schedule is always active.

Note you may not delete a schedule that is being used by any firewall rule or rule list.

EXAMPLES

```
create schedule my_schedule1 date-valid-start now date-valid-end 12/31/2016 daily-hour-start 8:00 daily-hour-end 17:00
```

Creates a new schedule which is active between 8am and 5pm every day until December 31, 2016.

```
list schedule>
```

```
security firewall schedule my_schedule1 {
```

```
  daily-hour-end 17:00
```

```
  daily-hour-start 8:00
```

```
  date-valid-end 2016-12-31:00:00:00
```

```
  date-valid-start 2012-12-12:08:40:01
```

```
}
```

```
security firewall schedule workHours {
```

```
  daily-hour-end 18:00
```

```
  daily-hour-start 8:00
```

```
  days-of-week { monday tuesday wednesday thursday friday }
```

```
}
```

Lists two user-configured schedules, including the one that you created above.

```
modify schedule my_schedule1 days-of-week { monday tuesday wednesday }
```

Modifies the schedule named "my_schedule1." This limits the schedule to running only on Mondays, Wednesdays, and Fridays.

OPTIONS

app-service

Associates this schedule with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers. The asm module has components for working with iApps.

description

Describes the schedule.

daily-hour-end

Specifies the time of day this schedule stops. This end hour must be after the daily-hour-start value. The default is 24:00 (midnight).

A schedule may not contain hours that go past midnight (24:00): for example, a daily-hour-start of 20:00 and daily-hour-end of 02:00 is not allowed. If you need to cover both the late hours and early hours of the day, please create two schedules.

daily-hour-start

Specifies the time of day this schedule starts. This start hour must be before the daily-hour-end value. The default is 0:00 (midnight at the start of the day).

date-valid-end

Specifies the final date for this schedule. The schedule stops firing as of this date. You may specify just the specific date, or a specific date and time for the schedule to end. The date must be after the date-valid-start value. The default is 19:14 1/18/2038 (the latest date expressible with a 32-bit integer).

date-valid-start

Specifies the start date for this schedule. The schedule does not fire before this date and time. You may specify just the specific date, or a specific date and time for the schedule to start. You must specify a date before the date-valid-end value. The default is midnight 1/1/1970 (Unix epoch).

days-of-week

Specifies which days of the week the schedule fires. You must specify at least one day of the week, and you cannot specify any day of the week more than once. The default is all seven days.

SEE ALSO

create, delete, edit, list, modify, net self, net route-domain, security firewall global-rules, security firewall management-ip-rules, security firewall rule-list, ltm virtual, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 security firewall schedule(1)

security firewall user-domain

NAME

user-domain - Configures a user-domain for use by firewall rules.

MODULE

security firewall

SYNTAX

CREATE/MODIFY

create user-domain [name]

modify user-domain [[name] | all]

options:

domain string

imap-service [add | delete | modify | replace-all-with] {

[ip]

[port]

[login]

[password]

}

app-service [name]

description [string]

edit user-domain [[name] | all]

options:

all-properties

non-default-properties

DISPLAY

list user-domain [[name] | all] [property]

show running-config user-domain [[name] | all] [property]

DELETE

delete user-domain [[name] | all]

RUN

run user-domain [[name] | all] clear-cache

DESCRIPTION

A user-domain object holds attributes to reach services that provide more information about a user. This information includes the domain for which the service is defined, the identity service that can be used to validate this user and the ifmap service that can be used to obtain more information about the user. User domains also have associated cache data which is utilized by user identity feature for performance improvements. You can use the user-domain component to define reusable configuration that is used to learn about more users used in the firewall rules or clear their associated cache data on demand.

EXAMPLES

```
create user-domain gladiators identity-server add { felix-legions { ip 1.1.1.1 login maximus password meridius port 10002 } }
```

Creates a new user-domain object that defines a new identity-server for domain gladiators with an identity service felix-legions

```
modify user-domain gladiators ifmap-server add { rome { ip 10.10.10.10 login marcus password aurelius port 10002 } }
```

Modifies the above user domain to add an ifmap server.

```
run user-domain grumpycat clear-cache
```

Clears the user identity cache associated with a specified user domain name.

```
run user-domain all clear-cache
```

Clears the user identity cache for all available user domains.

OPTIONS

domain

Specifies a domain for which the identity and ifmap services are defined.

app-service

Associates this user domain with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers. The asm module (see asm) has components for working with iApps.

description

User-defined description for this user domain.

clear-cache

Invokes clear-cache functionality for the given user-domain name.

SEE ALSO

edit, list, modify, net self, net route-domain, security firewall global-rules, security firewall management-ip-rules, security firewall rule-list, ltm virtual, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2014, 2016. All rights reserved.

BIG-IP 2016-03-14 security firewall user-domain(1)

security firewall user-list

NAME

user-list - Configures a user-list for use by firewall rules. A firewall rule can match a packet sourced from a particular user against one of the users or user-groups in a user list, and can take some action (such as ACCEPT or DROP) for a matching packet. An incoming packet's source IP address is matched in user identity database to get the user and group properties which are then used to perform the rule match.

MODULE

security firewall

SYNTAX

CREATE/MODIFY

```
create user-list [name]
```

```
modify user-list [[name] | all]
```

options:

```
app-service [name]
```

```
description [string]
```

```
user-groups [add | delete | modify | replace-all-with] {  
  [ [user group names...] ]  
}
```

```
users [add | delete | modify | replace-all-with] {
```

```
[ [user names...] ]  
}
```

```
edit user-list [[name] | all]  
options:  
  all-properties  
  non-default-properties
```

```
DISPLAY  
list user-list [[name] | all | [property]]
```

```
DELETE  
delete user-list [[name] | all]
```

DESCRIPTION

You can use the user-list component to define reusable lists of user or user-group names for various firewall rules. The network software compares a packet's source user (mapped by incoming source IP address) and group that user belong to, against users (or user-groups) in this list. You can assign a user list to the firewall rules in net self, net route-domain, security firewall global-rules, security firewall rule-list, and ltm virtual firewall rules.

EXAMPLES

```
create user-list u-list1 users add { olympus\xyz }
```

Creates a new user list named u-list1 with one user named xyz in domain olympus.

```
create user-list u-list2 user-groups add { olympus\eng }
```

Creates a new user list named u-list2 with one group named eng in domain olympus.

```
list user-list
```

Shows all the user lists configured in the system.

OPTIONS

```
app-service  
Associates this user list with a particular Application Service. An Application Service is a major component of an iApp, an advanced configuration tool for creating and maintaining similar applications on multiple servers. The asm module has components for working with iApps.
```

```
description  
Your description for the user list.
```

```
user-groups  
Specifies a list of user groups to compare against the groups a user belongs to (which is mapped from the source IP address).
```

```
users  
Specifies a list of users to compare against a packet's source user (which is mapped from the source IP address).
```

SEE ALSO

edit, list, modify, net self, net route-domain, security firewall address-list, security firewall rule-list, security firewall global-rules, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 security firewall user-list(1)

security firewall uuid-default-autogenerate

NAME
uuid-default-autogenerate - Configures system wide firewall rule uuid auto generate mode.

MODULE
security firewall

SYNTAX
MODIFY
modify uuid-default-autogenerate mode [disabled | enabled]

DISPLAY
list uuid-default-autogenerate mode
show uuid-default-autogenerate mode

DESCRIPTION

You can use the `uuid-default-autogenerate` component to change the behavior of how the firewall rule `uuid` get auto generated. By default the system will automatically disable firewall rule `uuid` auto-generating. You can change the behavior to manually trigger the rule `uuid` auto-generating by setting the mode to enabled.

EXAMPLES

modify security firewall `uuid-default-autogenerate` mode enabled

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014, 2017. All rights reserved.

BIG-IP 2018-10-20 security firewall `uuid-default-autogenerate(1)`

security flowspec-route-injector flowspec-advertised-route-info

NAME

`flowspec-advertised-route-info` - Display FlowSpec routes that are currently being advertised.

MODULE

security `flowspec-route-injector`

SYNTAX

Display `flowspec-advertised-route-info` component within the `security flowspec-route-injector` module using the syntax in the following section.

DISPLAY

show `flowspec-advertised-route-info route-domain-id`

DESCRIPTION

Show `flowspec-advertised-route-info` displays FlowSpec routes that are currently being advertised per route domain instance.

OPTIONS

For information about the options that you can use with the command `show`, see `help show`.

SEE ALSO

`show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP security `flowspec-route-injector flowspec-advertised-route-info(1)`

security flowspec-route-injector profile

NAME

`profile` - Configures a Security FlowSpec Route Injector profile

MODULE

security `flowspec-route-injector`

SYNTAX

Manage profile component within the `security flowspec-route-injector` module using the syntax shown in the following sections.

CREATE/MODIFY

create profile [name]

modify profile [name]

options:

app-service [[string] | none]

description [string]

max-flowspec-routes-limit [integer]

neighbor [add | delete | modify | none | replace-all-with] {

```

    [IP Address] {
adj-out [disabled | enabled]
bgp-multiple-instance [disabled | enabled]
extended-asn-cap [disabled | enabled]
graceful-restart [disabled | enabled]
graceful-restart-time [integer]
hold-time [integer]
local-address [IP Address]
local-as [integer]
remote-as [integer]
router-id [IPv4 Address]
    }
}
rules [[add | delete | modify | none | replace-all-with] {
    [name] {
action {
    dscp-value [integer]
    next-hop [IP Address]
    rate-limit [integer]
    asn-community [string]
    type [drop | redirect | rate-limit | qos]
}
alias [string]
app-service [string]
advertisement-ttl-from-now [integer]
description [string]
remove-config-upon-expiry [bool]
match {
    destination-address [IP Address]
    destination-ports [list of ports / port-ranges]
    dscp-values [list of integers]
    icmp-codes [list of integers]
    icmp-types [list of integers]
    ip-fragments [list of integers]
    ip-protocols [list of protocols]
    packet-lengths [list of integers / integer-ranges]
    ports [list of ports / port-ranges]
    source-address [IP Address]
    source-ports [list of ports / port-ranges]
    tcp-flags {
        bitwise-and-fields [list of integers]
        bitwise-or-fields [list of integers]
    }
}
}
}
route-domain [name]
peer-group {
adj-out [disabled | enabled]
bgp-multiple-instance [disabled | enabled]
extended-asn-cap [disabled | enabled]
graceful-restart [disabled | enabled]
graceful-restart-time [integer]
hold-time [integer]
local-address [IP Address]
local-as [integer]
remote-as [integer]
router-id [IPv4 Address]
}
security-log-profile [string]

edit profile
options:
    all-properties
    non-default-properties

DISPLAY
list profile
show running-config profile

```

DESCRIPTION

profile component under security flowspec-route-injector is used to manage a Security FlowSpec Route Injector profile (unique per route domain instance). Security FlowSpec route injector profile is used by AFM/DHD module to advertise routes based on Source/Destination IP, Source/Destination Port, Protocol etc. for blackholing and scrubbing use cases using BGP FlowSpec mechanism (RFC 5575).

EXAMPLES

```

create profile p1
neighbor add {
    10.128.10.128 {
local-address 10.128.10.169
    }
}
peer-group {
local-as 60000
remote-as 60000
router-id 1.1.1.1
}

```

```
}  
route-domain 0 }
```

Create a security flowspec-route-injector profile p1 for route-domain 0 and add 1 peer neighbor 10.128.10.128. Common attributes that are shared by all neighbors in the profile (unless overridden) are defined using peer-group settings.

```
modify profile p1 peer-group { graceful-restart enabled graceful-restart-time 120 }
```

Modify profile p1 and update graceful-restart and graceful-restart-time peer-group attributes.

```
list policy
```

Displays the current list of configured security flowspec-route-injector profiles.

OPTIONS

description

User defined description.

advertisement-ttl-from-now

Specifies the duration (in minutes) after which FlowSpec should be withdrawn. The default is 5 minutes. If it is 0, it would be allowed for user to immediately expire the rule (and withdraw from upstream routers).

This is user write-only configuration. It is used for system to calculate expiry time of the rule. It is mutual exclusive with expiry-time.

max-flowspec-routes-limit

Specifies the maximum number of FlowSpec routes that can be advertised simultaneously per FlowSpec profile (or route domain) instance. Minimum allowed value is 100, Maximum allowed value is 10,000 (which is default value too).

neighbor

Add, modify, delete BGP peer neighbor configuration. Each neighbor is uniquely identified / configured using IP Address as the name.

description

User defined description.

adj-out

Enable/Disable BGP adj-rib-out feature. Default is enabled.

bgp-multiple-instance

Enable/Disable BGP multiple instance capability. Default is disabled.

extended-asn-cap

Enable/Disable Extended ASN capability (i.e. send 4-byte ASN). Default is enabled.

graceful-restart

Enable/Disable graceful restart capability. Default is disabled.

graceful-restart-time

Specifies graceful restart time (max time needed for Neighbor(s) to restart).

hold-time

Specifies the hold time (max time that can elapse between messages from peer). Default is 90 seconds.

local-address

Specifies the Local Address (on BigIP) to be used for initiating BGP connection(s) with peers.

local-as

Specifies the BGP Local AS number.

remote-as

Specifies the BGP Remote AS number.

router-id

Specifies the BGP Router ID to be used in BGP OPEN message when initiating BGP connection with peers. Router ID is an IPv4 address.

route-domain

Specifies name of the route domain to be used by the Security FlowSpec Route Injector profile. This is required field at the time of profile creation and is non-mutable after policy creation.

rules

Specifies configuration of rules that can be advertised per FlowSpec profile.

action

Specifies BGP FlowSpec Advertisement Action configuration.

dscp-value

Specifies the BGP FlowSpec DSCP value for advertisement qos action. The default is 0. The valid range is 0 ~ 63 inclusive.

next-hop

Specifies BGP FlowSpec redirection next hop address

rate-limit

Specifies the BGP FlowSpec rate limit (bytes/sec) for advertisement rate limiting action.

asn-community

Specifies the BGP Extended Community value (in the format - AA:NNN, where AA is 16-bit number and NNN is 32-bit number) for redirect-to-VRF support when BGP Flowspec advertisement action is redirect.

type Specifies the BGP FlowSpec Advertisement Action type for this FlowSpec Route Injector profile.

The default is redirect.

alias

Specifies the alias name of this rule.

app-service

The application service that the object belongs to.

creation-time

The time when this rule is created. This is not user configurable field.

description

User defined description.

expiry-time

The time when this rule is going to be expired.

This field is mutual exclusive with advertisement-ttl-from-now. If user specifies advertisement-ttl-from-now, expiry-time will be calculated from it.

last-modified-time

The time when this rule is modified. This is not user configurable field.

remove-config-upon-expiry

Specifies whether or not this rule needs to be automatically removed when reaching expiry time. The default is true. If it is set to false, user needs to manually remove this rule as it is needed.

The maximum allowed expired rules per profile in database is defined by DB variable, flowspec.max.expired_and_saved_rules (min = 0, max = 1000, default = 100).

match

Specifies BGP FlowSpec matching criteria configuration.

destination-address

Specifies the destination address/prefix to match in packets.

destination-ports

Specifies a list of ports that matches destination TCP/UDP ports in packets.

This destination-ports configuration is mutual exclusive with ports field.

dscp-values

Specifies a list of DSCP values to match in packets. The valid range for each of DSCP value in the list must be within 0 ~ 63 inclusive.

icmp-codes

Specifies a list of ICMP codes to match in packets.

icmp-types

Specifies a list of ICMP types to match in packets.

ip-fragments

Specifies a list of IP fragments to match in packets.

ip-protocols

Specifies a set of protocol values that are used to match the IP protocol value byte in IP packets. The valid protocols are ICMP, TCP, UDP, and SCTP. If port object is specified, the valid protocols are TCP, UDP, and SCTP.

packet-lengths

Specifies a list of packet lengths (singleton or a range) to match. Packet Length includes L3 (header) size in addition to payload length.

ports

Specifies a list of ports that matches source OR destination TCP/UDP ports in packets.

This ports configuration is mutual exclusive with destination-ports and source-ports.

source-address

Specifies the source address/prefix to match in packets.

source-ports

Specifies a list of ports that matches source TCP/UDP ports in packets.

This source-ports configuration is mutual exclusive with ports field.

tcp-flags

Specifies lists of TCP flags to match in packets."

`bitwise-and-fields`

Specifies a bitwise AND list of TCP flags to match in packets."

`bitwise-or-fields`

Specifies a bitwise OR list of TCP flags to match in packets."

`peer-group`

Specifies peer group settings that are inherited by each neighbor unless overridden specifically for that neighbor.

`adj-out`

Enable/Disable BGP adj-rib-out feature. Default is enabled.

`bgp-multiple-instance`

Enable/Disable BGP multiple instance capability. Default is disabled.

`extended-asn-cap`

Enable/Disable Extended ASN capability (i.e. send 4-byte ASN). Default is enabled.

`graceful-restart`

Enable/Disable graceful restart capability. Default is disabled.

`graceful-restart-time`

Specifies graceful restart time (max time needed for Neighbor(s) to restart).

`hold-time`

Specifies the hold time (max time that can elapse between messages from peer). Default is 90 seconds.

`local-address`

Specifies the Local Address (on BigIP) to be used for initiating BGP connection(s) with peers.

`local-as`

Specifies the BGP Local AS number.

`remote-as`

Specifies the BGP Remote AS number.

`router-id`

Specifies the BGP Router ID to be used in BGP OPEN message when initiating BGP connection with peers. Router ID is an IPv4 address.

`security-log-profile`

Specifies log publisher name used for this FlowSpec Route Injector profile.

SEE ALSO

`create`, `edit`, `list`, `modify`, `security`, `security scrubber`, `security scrubber profile`, `security blacklist-publisher profile`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015, 2017. All rights reserved.

BIG-IP 2018-12-28 security flowspec-route-injector profile(1)

security http file-type

NAME

`file-type` - Lists the available file types that can be used in the context of HTTP Protocol Security.

MODULE

`security http`

SYNTAX

Retrieve the list of the `file-type` values using the syntax shown in the following sections.

DISPLAY

`list file-type`

`list file-type [[[name] | [glob] | [regex]] ...]`

options:

`all`

`app-service`

`one-line`

DESCRIPTION

Use this command to display the possible values of the file-type object to be used in the context of HTTP Protocol Security. These possible values include predefined and user-defined file types that you can select to have the security profiles allow or disallow.

EXAMPLES

```
list file-type
```

Displays all the file types supported by HTTP Protocol Security.

OPTIONS

```
app-service
```

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

SEE ALSO

glob, list, regex, security http profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-06-12 security http file-type(1)

security http mandatory-header

NAME

mandatory-header - Lists the available mandatory headers that can be used in the context of HTTP Protocol Security.

MODULE

security http

SYNTAX

Retrieve the list of the mandatory-header values using the syntax shown in the following sections.

DISPLAY

```
list mandatory-header
```

```
list mandatory-header [ [ [name] | [glob] | [regex] ] ... ]
```

options:

```
all
app-service
one-line
```

DESCRIPTION

Use this command to display the possible values of the mandatory-header object to be used in the context of HTTP Protocol Security. These possible values include predefined and user-defined HTTP headers that you can select to be required by the security profiles.

EXAMPLES

```
list mandatory-header
```

Displays all the mandatory headers supported by HTTP Protocol Security.

OPTIONS

```
app-service
```

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

SEE ALSO

glob, list, regex, security http profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-06-12 security http mandatory-header(1)

security http profile

NAME

profile - Configures an HTTP security profile.

MODULE

security http

SYNTAX

Configure the profile component within the security http module using the syntax shown in the following sections.

CREATE/MODIFY

```
create profile [name]
```

```
modify profile [name]
```

```
options:
```

```
  app-service [[string] | none]
  [case-sensitive | case-insensitive]
```

```
  defaults-from [[name] | none]
```

```
  description [[string] | none]
```

```
  evasion-techniques {
```

```
    options:
```

```
  alarm [disabled | enabled]
```

```
  block [disabled | enabled]
```

```
  }
```

```
  file-types {
```

```
    options:
```

```
  alarm [disabled | enabled]
```

```
  [allowed | disallowed]
```

```
  block [disabled | enabled]
```

```
  values [add | delete | none | replace-all-with] { [string] ... }
```

```
  }
```

```
  http-rfc {
```

```
    options:
```

```
  alarm [disabled | enabled]
```

```
  bad-host-header [disabled | enabled]
```

```
  bad-version [disabled | enabled]
```

```
  block [disabled | enabled]
```

```
  body-in-get-head [disabled | enabled]
```

```
  chunked-with-content-length [disabled | enabled]
```

```
  content-length-is-positive [disabled | enabled]
```

```
  header-name-without-value [disabled | enabled]
```

```
  high-ascii-in-headers [disabled | enabled]
```

```
  host-header-is-ip [disabled | enabled]
```

```
  maximum-headers [[integer] | disabled]
```

```
  null-in-body [disabled | enabled]
```

```
  null-in-headers [disabled | enabled]
```

```
  post-with-zero-length [disabled | enabled]
```

```
  several-content-length [disabled | enabled]
```

```
  unparseable-content [disabled | enabled]
```

```
  }
```

```
  mandatory-headers {
```

```
    options:
```

```
  alarm [disabled | enabled]
```

```
  block [disabled | enabled]
```

```
  values [add | delete | none | replace-all-with] { [string] ... }
```

```
  }
```

```
  maximum-length {
```

```
    options:
```

```
  alarm [disabled | enabled]
```

```
  block [disabled | enabled]
```

```
  post-data [[integer] | any]
```

```
  query-string [[integer] | any]
```

```
  request [[integer] | any]
```

```
  uri [[integer] | any]
```

```
  }
```

```
  methods {
```

```
    options:
```

```
  alarm [disabled | enabled]
```

```
  block [disabled | enabled]
```

```
  values [add | delete | none | replace-all-with] { [string] ... }
```

```
  }
```

```
  response {
```

```
    options:
```

```
  body [[string] | none]
```

```
  headers [[new line separated headers] | none]
```

```
  type [custom | default | redirect | soap-fault]
```

```
  uri [[string] | none]
```

```
  }
```

```
edit profile [ [name] | [glob] | [regex] ] ... ]
```

```
options:
```

```
  all-properties
```

non-default-properties

DISPLAY

list profile

list profile [[name] | [glob] | [regex]] ...]

show running-config profile

show running-config profile [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

partition

recursive

DELETE

delete profile [name]

DESCRIPTION

You can use the profile component to create, modify, display, or delete an HTTP security profile for use with HTTP Protocol Security functionality.

EXAMPLES

```
create http my_http_profile defaults-from http_security
```

Creates a custom HTTP security named my_http_profile that inherits its settings from the system default HTTP security profile.

```
list profile
```

Displays the properties of all HTTP security profiles.

OPTIONS

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

[case-sensitive | case-insensitive]

Specifies whether the security profile treats file types as case sensitive, or not. The default value is case-sensitive. Note: If you create a profile, you can use either property, thereafter it becomes read only. If the security profile is case insensitive, the system stores file types in lowercase in the security profile configuration.

defaults-from

Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified. The default value is none.

description

User defined description.

evasion-techniques

Specifies what action the system takes when it detects an evasion technique. Evasion techniques are methods used by attackers to avoid detection of their attack. You can configure the following options for evasion technique checks:

alarm

Specifies, when enabled, that the system logs the request data and displays it in the Protocol Security Statistics screen whenever the system detects an evasion technique. The default value is enabled.

block

Specifies, when enabled, that the system stops requests whenever the system detects an evasion technique. The default value is disabled.

file-types

Specifies which file types the security profile considers legal, and specifies what action the system takes when it detects a request for an illegal file type. You can configure the following options for file types:

alarm

Specifies, when enabled, that the system logs the request data and displays it on the Protocol Security Statistics screen whenever the system detects a request for an illegal file type. The default value is enabled.

[allowed | disallowed]

Indicates whether the values property lists file types that the security profile permits or prohibits. Note: For each security profile you may define either allowed file types or disallowed file types.

block

Specifies, when enabled, that the system stops requests for an illegal file type. The default value is disabled.

values

Adds, deletes, or replaces a set of file types considered either legal or illegal by the security profile. You can either select an available file-type or add a new one.

`glob` Displays the items that match the `glob` expression. See help `glob` for a description of `glob` expression syntax.

`http-rtc`

Specifies which validations the system should check and what action the system takes when it detects a request that is not formatted properly. You can configure the following options for HTTP protocol checks:

`alarm`

Specifies, when enabled, that the system logs the request data and displays it in the Protocol Security Statistics screen whenever a request fails one of the enabled HTTP protocol checks. The default value is enabled.

`bad-host-header`

Specifies, when enabled, that the system inspects requests to see whether they contain a non RFC compliant header value. The default value is enabled.

`bad-version`

Specifies, when enabled, that the system inspects requests to see whether they request information from a client using an HTTP protocol version 1.0 or higher. The default value is enabled.

`block`

Specifies, when enabled, that the system stops requests whenever the system detects an evasion technique. The default value is disabled.

`body-in-get-head`

Specifies, when enabled, that the system examines requests that use the HEAD or GET methods to see whether the requests contain data in their bodies, which is considered illegal. The default value is disabled.

`chunked-with-content-length`

Specifies, when enabled, that the system examines chunked requests for a content-length header, which is not permitted. The default value is enabled.

`content-length-is-positive`

Specifies, when enabled, that the system examines requests to see whether their content length value is greater than zero. The default value is enabled.

`header-name-without-value`

Specifies, when enabled, that the system checks requests for valueless header names, which are considered illegal. The default value is enabled.

`high-ascii-in-headers`

Specifies, when enabled, that the system inspects request headers for ASCII characters greater than 127, which are not permitted. The default value is disabled.

`host-header-is-ip`

Specifies, when enabled, that the system verifies that the request's host header value is not an IP address. The default value is disabled.

`maximum-headers`

Specifies whether the system compares the number of headers in the requests against the maximum number, and if so, how many headers are allowed. The default value is a maximum of 20 headers.

`null-in-body`

Specifies, when enabled, that the system inspects request bodies to see whether they contain a Null character, which is not allowed. The default value is disabled.

`null-in-headers`

Specifies, when enabled, that the system inspects request headers to see whether they contain a Null character, which is not allowed. The default value is enabled.

`post-with-zero-length`

Specifies, when enabled, that the system examines POST method requests for no content-length header, and for a content length of 0. The default value is disabled.

`several-content-length`

Specifies, when enabled, that the system examines each request to see whether it has more than one content-length header, which is considered illegal. The default value is enabled.

`unparsable-content`

Specifies, when enabled, that the system examines requests for content that the system cannot parse, which is not permitted. The default value is enabled.

`mandatory-headers`

Specifies which headers must appear in requests, and specifies what action the system takes when it detects a request without a mandatory header. You can configure the following options for mandatory headers:

`alarm`

Specifies, when enabled, that the system logs the request data and displays it on the Protocol Security Statistics screen whenever a request does not include a mandatory header. The default value is enabled.

`block`

Specifies, when enabled, that the system stops requests that do not include a mandatory header. The default value is disabled.

values

Adds, deletes, or replaces a set of headers that must appear in requests to be considered legal by the security profile. You can either select an available mandatory-header or add a new one. Note: The system stores mandatory headers in lowercase in the security profile configuration, regardless of whether it is case sensitive or not.

maximum-length

Specifies the default maximum length settings that the security profile considers legal, and specifies what action the system should take when it detects a request using an illegal length. You can configure the following options for length checks:

alarm

Specifies, when enabled, that the system logs the request data and displays it on the Protocol Security Statistics screen whenever a request fails one of the length checks. The default value is enabled.

block

Specifies, when enabled, that the system stops requests that fail one of the length checks. The default value is disabled.

post-data

Indicates whether there is a maximum acceptable length, in bytes, for the POST data portion of a request, and if so, specifies it. The default value is any (no restriction).

query-string

Indicates whether there is a maximum acceptable length, in bytes, for the query string portion of a request, and if so, specifies it. The default value is 1024 bytes.

request

Indicates whether there is a maximum acceptable length, in bytes, of a request, and if so, specifies it. The default value is any (no restriction).

uri Indicates whether there is a maximum acceptable length, in bytes, for a URL, and if so, specifies it. The default value is 1024 bytes.

methods

Specifies which HTTP methods the security profile considers legal, and specifies what action the system takes when it detects a request using an illegal method. You can configure the following options for methods:

alarm

Specifies, when enabled, that the system logs the request data and displays it on the Protocol Security Statistics screen whenever a request uses an illegal method. The default value is enabled.

block

Specifies, when enabled, that the system stops requests that use an illegal method. The default value is disabled.

values

Adds, deletes, or replaces a set of HTTP methods considered legal by the security profile. You can either select an available asm http-method or add a new one. Note: HTTP methods are case sensitive even if the security profile is case insensitive.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

response

Specifies information to display when the security profile blocks a client request. You can configure the following options for blocking page:

body Specifies the HTML code the system sends to the client in response to an illegal blocked request. Only if the response type is custom, you can edit this text.

headers

Specifies the set of response headers that the system sends to the client in response to an illegal blocked request. Only if the response type is custom, you can edit this text. Separate each header with a new line (Ctrl-V followed by Ctrl-J).

type Specifies which content, or URL, the system sends to the client in response to an illegal blocked request.

custom

Specifies a modified response text. You can edit the response header and HTML code in the properties headers and body.

default

Specifies the system-supplied response text written in HTML. You cannot edit that text. This is the default value.

redirect

Specifies that the system redirects the user to a specific web page instead of viewing a blocking page. You can edit the redirect web page in the url property.

soap-fault

Specifies the system-supplied response written in SOAP fault message structure. You cannot edit that text. Use this type when a SOAP request is blocked due to an XML related violation.

url Specifies the particular URL to which the system redirects the user. Only if the response type is redirect, you can edit this text. The web page should include a full URL path, for example, <http://www.myredirectpage.com>.

SEE ALSO

asm http-method, create, delete, edit, glob, list, ltm virtual, modify, regex, security, security http, security http file-type, security http mandatory-header, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2017-05-24 security http profile(1)

security ip-intelligence blacklist-category

NAME

blacklist-category - Global list of ip-intelligence blacklist categories. These ip-intelligence blacklist categories are used to configure ip-intelligence policies.

MODULE

security ip-intelligence

SYNTAX

Configure the blacklist-category component within the security ip-intelligence module using the syntax shown in the following sections.

CREATE/MODIFY

create blacklist-category [name]

modify blacklist-category [[name] | all]

options:

app-service [name]

description [string]

bl-match-direction [destination | source | source-and-destination]

edit blacklist-category

options:

all-properties

non-default-properties

one-line

partition

recursive

DISPLAY

list blacklist-category

show running-config blacklist-category

options:

all-properties

non-default-properties

one-line

partition

recursive

DESCRIPTION

You can use the blacklist-category component to configure a shareable and reusable blacklist category which can be configured with specific enforcement and logging settings under ip-intelligence policies.

EXAMPLES

modify blacklist-category Malware description "A variety of forms of hostile or intrusive software."

Modifies the blacklist-category description.

list blacklist-category

Displays the current list of blacklist categories.

OPTIONS

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description
User defined description.

bl-match-direction
Indicates whether to match source IPs, destination IPs, or both.

partition
Displays the administrative partition within which the component resides.

SEE ALSO

create, edit, list, modify, security ip-intelligence feed-list, security ip-intelligence policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 security ip-intelligence blacklist-category(1)

security ip-intelligence feed-list

NAME

feed-list - Configures a feed-list for use by firewall. A feed-list is a list of URL feeds from where files are downloaded and the contents (IP-address prefixes) are compared against the source-IP address and/or destination-IP address in an IP packet by DWBL (Dynamic White/Black lists) by IP-Intelligence.

MODULE

security ip-intelligence

SYNTAX

Configure the feed-list component within the security ip-intelligence module using the syntax in the following sections.

CREATE/MODIFY

```
create feed-list [name]
modify feed-list [[name] | all]
options:
feeds [add | delete | modify | replace-all-with] {
  name [string] {
    options:
default-blacklist-category [string]
default-list-type [whitelist | blacklist]
poll {
  interval [integer]
  user [string]
  url [string]
  password [string]
}
}
}
app-service [name]
description [string]
```

```
edit feed-list [[name] | all]
```

```
options:
  all-properties
  non-default-properties
```

```
load feed-list [[name] | all] feeds { name [string] }
```

DISPLAY

```
list feed-list [[name] | all | [property]]
show running-config feed-list [[name] | all | [property]]
```

```
options:
  all-properties
  non-default-properties
  one-line
  partition
  recursive
```

DELETE

```
delete feed-list [[name] | all]
```

DESCRIPTION

You can use the feed-list component to define reusable lists of feeds. You can use a feed list in a security ip-intelligence policy. A policy compares all of the addresses in the list (downloaded from a file at the specified url) to either the source or destination IP in the packet, depending on how you apply the list. If there is a match, the ip-intelligence policy takes an action, such as accepting or dropping the packet.

EXAMPLES

```
create feed-list alist1 feeds add { poll { url http://f5.com/bl.txt }
```

Creates a new feed list, "alist1," with IPv4/IPv6 addresses in the file downloaded from the specified url.

```
modify feed-list alist1 feeds modify { description "DWBL file from f5.com" }
```

Modifies the above feed list with a description.

```
modify feed-list alist1 feeds modify { poll { url https://f5.com/bl.txt }
```

Modifies the same feed by changing the protocol.

```
list feed-list alist1
security ip-intelligence feed-list alist1 {
  feeds {
    url2 {
      poll {
        url https://f5.com/bl.txt
        user user1
        password user1_pwd
      }
    }
    description "DWBL file from f5.com"
  }
}
```

Shows the modified feed list.

```
load feed-list alist1 alist2 feeds { feed1 feed2 }
```

Immediately downloads and updates feeds feed1 and feed2 of feed lists alist1 and alist2.

OPTIONS

feeds

Adds, deletes, or replaces feeds. You can configure the following options for a feed:

name Specifies a name for a feed. This option is required for the operations create, delete, modify, and replace-all-with.

add Creates a new feed list.

delete

Deletes the feed list that you specify next, in curly braces ({}).

modify

Makes it possible to replace the optional description(s) for the feed list.

replace-all-with

Replaces the current set of feed list with the a new one that you specify next, in curly braces ({}).

default-list-type

Specifies a default type for this specific entry whether it is a blacklist or whitelist

whitelist

Specifies that this entry is a whitelist.

blacklist

Specifies that this entry is a blacklist.

default-blacklist-category

Default blacklist category type for all blacklist entries that do not have a corresponding category string (eg. Botnet, Spyware, Malware)

poll You can configure the following options under this:

interval

Specifies the frequency at which the url needs to be polled.

user Specifies the user which is used when downloading the url.

url Specifies the URL from where the white/black list will be downloaded. Note: Route domains are not supported when specifying the url.

password

Password for the user.

default-list-type

Specifies a default type for this specific entry whether it is a blacklist or whitelist

whitelist

Specifies that this entry is a whitelist.

blacklist

Specifies that this entry is a blacklist.

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

default-blacklist-category

Default blacklist category type for all blacklist entries that do not have a corresponding category string (eg. Botnet, Spyware, Malware)

description

User defined description for this feed list.

partition

Displays the administrative partition within which the component resides.

SEE ALSO

edit, list, modify, net self, net route-domain, security ip-intelligence global-policy, security ip-intelligence, ltm virtual, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 security ip-intelligence feed-list(1)

security ip-intelligence global-policy

NAME

global-policy - Configures the global ip-intelligence policy. These ip-intelligence policy contents/filters are applied to all packets except those going through the management interface. They are applied first, before any firewall rules for the packet's virtual server, route domain.

MODULE

security ip-intelligence

SYNTAX

Modify the global-policy component within the security ip-intelligence module using the syntax shown in the following sections.

MODIFY

modify global-policy

options:

app-service [name]

description [string]

ip-intelligence-policy [[policy_name] | none]

edit global-policy

options:

all-properties

non-default-properties

one-line

partition

recursive

reset-stats global-policy

options:

ip-intelligence-categories

DISPLAY

list global-policy

show running-config global-policy

options:

all-properties

non-default-properties

one-line

partition

recursive

show global-policy

options:

ip-intelligence-categories

DESCRIPTION

You can use the global-policy component to configure a shareable and reusable set of network firewall DWBL (Dynamic White/Black lists) which can be enforced globally at the system level and the enforcement happens before the route-domain or virtual server level.

EXAMPLES

modify global-policy policy pol1

Modifies the global-policy with policy pol1.

list global-policy

Displays the current list of global-policy contents.

OPTIONS

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description

User defined description.

policy

Specifies an existing policy. policy contents are enforced at a global level.

partition

Displays the administrative partition within which the component resides.

ip-intelligence-categories

Used to show/ reset statistics on IP intelligence white/ black lists categories.

SEE ALSO

create, edit, list, modify, security ip-intelligence feed-list, security ip-intelligence policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 security ip-intelligence global-policy(1)

security ip-intelligence info

NAME

info - Query IP Intelligence information about specified IP Address, GEO location or FQDN.

MODULE

security ip-intelligence

SYNTAX

Use the info component within the security ip-intelligence module to query IP Intelligence information about the IP address, GEO location or FQDN using the following syntax.

DISPLAY

```
show info address [IP address]
fqdn [FQDN name]
geo [GEO_country[:GEO_region]]
options:
virtual-server [name]
route-domain [name]
```

DESCRIPTION

You can use the info component to query IP Intelligence information about an IP Address, a GEO location, or a FQDN using IP Intelligence policy attached to the selected context. Global context is the default when the command is used without options. To select the virtual server or the route domain context use virtual-server or route-domain option. Only one option can be used. If the specified parameter (IP Address, GEO location or FQDN) is listed in any of the DWBL (Dynamic White/Black List) feeds used by the selected IP Intelligence policy the query shows the list of categories and policy action (drop or allow) for the address. If the policy is configured to query legacy IP Reputation database, that information is also used in the query.

EXAMPLES

```
show info address 10.123.1.12
```

Query IP Intelligence information for IP address 10.123.1.12 using global IP Intelligence policy.

```
show info address 10.123.1.12 virtual-server /Common/vs
```

Query IP Intelligence information for IP address 10.123.1.12 using IP Intelligence policy configured for virtual server /Common/vs.

```
show info address fqdn { f5.com }
```

Query IP Intelligence information for FQDN f5.com using global IP Intelligence policy.

```
show info address geo { US }
```

Query IP Intelligence information for United States using global IP Intelligence policy.

```
show info address geo { US:California }
```

Query IP Intelligence information for California, United States using global IP Intelligence policy.

OPTIONS

virtual-server

Specifies the name of the virtual server configured with IP Intelligence policy to use in the query. This option cannot be used with the route-domain option.

route-domain

Specifies the name of the route domain configured with IP Intelligence policy to use in the query. This option cannot be used with the virtual-server option.

SEE ALSO

security ip-intelligence feed-list, security ip-intelligence policy, tmsh, ltm virtual, net route-domain

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2018. All rights reserved.

BIG-IP 2018-09-20 security ip-intelligence info(1)

security ip-intelligence policy

NAME

policy - Configures an ip-intelligence policy. It's comprised of three logical groups of settings: list of feed lists, enforcement and logging settings per blacklist category, and default enforcement and logging settings for blacklist categories.

MODULE

security ip-intelligence

SYNTAX

Configure the policy component within the security ip-intelligence module using the syntax in the following sections.

CREATE/MODIFY

```
create policy [name]
```

```
modify policy [name]
```

options:

```
app-service [name]
```

```
description [string]
```

```
blacklist-categories [add | default | delete | replace-all-with] {
```

```
  [name] {
```

```
    action [accept | drop | use-policy-setting]
```

```
    app-service none
```

```
    description none
```

```
    log-blacklist-hit-only [no | yes | use-policy-setting]
```

```
    log-blacklist-whitelist-hit [no | yes | use-policy-setting]
```

```
    match-direction-override [match-destination | match-source | match-source-and-destination]
```

```
  }
```

```
}
```

```
feed-lists [add | default | delete | replace-all-with] { [name] }
```

```
default-action [accept | drop]
```

```
default-log-blacklist-hit-only [ no | yes ]
```

```
default-log-blacklist-whitelist-hit [ no | yes ]
```

```
edit policy
```

options:

```
all-properties
```

```
non-default-properties
```

DISPLAY

```
list policy [ [ [name] | [glob] | [regex] ] ... ]
show running-config policy
show running-config policy [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
  partition
  recursive
```

DESCRIPTION

You can use the policy component to configure a shareable and reusable enforcement and logging settings on Dynamic White/Black lists of IPs coming from downloaded feeds. The policy can then be enforced on a number of configuration objects of the following types: ltm virtual, security ip-intelligence global-policy, net route-domain.

EXAMPLES

```
create policy pol1 {
  blacklist-categories add {
    Spyware {
      action use-policy-setting
      app-service none
      description none
      log-blacklist-hit-only use-policy-setting
      log-blacklist-whitelist-hit yes
    }
  }
  feed-lists add { alist1 alist2 }
  default-action drop
  default-log-blacklist-hit-only yes
  default-log-blacklist-whitelist-hit no
  description none
  feed-lists none
  partition Common }
```

Creates a policy pol1 with feeds from alist1 and alist2 feed lists, specific enforcement and logging settings for Spyware blacklist category and policy default settings for other categories.

```
modify policy pol1 { feed-lists delete { alist2 } }
```

Removes the feed-list alist2 from the policy pol1.

```
list policy
```

Displays the current list of ip-intelligence policies contents.

OPTIONS

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description

User defined description.

partition

Displays the administrative partition within which the component resides.

blacklist-categories

Adds, deletes, or replaces blacklist categories.

action

Specifies what enforcement action will be applied if the packet is categorized with this blacklist category. If the packet is categorized with more than one blacklists the most restrictive action will be applied.

log-blacklist-hit-only

Specifies if a log message will be generated if the packet is categorized with this blacklist and the packet's IP listed in no whitelists.

match-direction-override

Overrides the current IP match direction setting for a category. If this value has not been overridden, it will be set to the value of the parent category's bl-match-direction at the time that the category was added to the policy.

log-blacklist-whitelist-hit

Specifies if a log message will be generated if the packet is categorized with this blacklist and the packet's IP is listed in a whitelist.

feed-lists

Adds, deletes, or replaces a feed list. Specifies a list of feed lists (see security ip-intelligence feed-list) against which the packet will be compared.

default-action

Specifies a default enforcement action which will be performed on the matched packet unless an implicit action specified for one of the blacklist categories the packet's IP is categorized with. If the packet's

IP is listed in a white list the action is always accept.

default-log-blacklist-hit-only

Specifies a default blacklist hit only logging action which will be performed on the matched packet unless an implicit action specified for one of the blacklist categories the packet's IP is categorized with.

default-log-blacklist-whitelist-hit

Specifies a default blacklist and whitelist hit logging action which will be performed on the matched packet unless an implicit action specified for one of the blacklist categories the packet's IP is categorized with.

SEE ALSO

create, edit, list, modify, security ip-intelligence feed-list, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 security ip-intelligence policy(1)

security log antifraud-storage-field

NAME

antifraud-storage-field - Lists the available storage format fields that can be used in the context of Anti-Fraud Security Logging.

MODULE

security log

SYNTAX

Retrieve the list of the antifraud-storage-field values using the syntax shown in the following sections.

DISPLAY

list antifraud-storage-field

list antifraud-storage-field [[[name] | [glob] | [regex]] ...]

options:

all
app-service
id
one-line
support-events

DESCRIPTION

Use this command to display the possible values of the antifraud-storage-field object to be used in the context of Anti-Fraud Security Logging. These possible values are predefined traffic items available for the server to log.

EXAMPLES

list antifraud-storage-field

Displays all the storage fields supported by Anti-Fraud Security Logging.

OPTIONS

app-service

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

id Displays an order ID of the field. This ID should be used to specify that this antifraud-storage-field should be URL-encoded when logging Anti-Fraud events.

support-events

Displays a bit mask representing the Anti-Fraud events that are supported by this antifraud-storage-field.

SEE ALSO

glob, list, regex, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012, 2018. All rights reserved.

security log network-storage-field

NAME

network-storage-field - Lists the available storage format fields that can be used in the context of Network Security Logging.

MODULE

security log

SYNTAX

Retrieve the list of the network-storage-field values using the syntax shown in the following sections.

DISPLAY

```
list network-storage-field
list network-storage-field [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  one-line
  app-service
```

DESCRIPTION

Use this command to display the possible values of the network-storage-field object to be used in the context of Network Security Logging. These possible values are predefined traffic items available for the server to log in context of Network event logging (for example, ACL events, TCP Open/Close, TCP/IP error events).

EXAMPLES

```
list network-storage-field
```

Displays all the storage fields supported by Network Security Logging.

OPTIONS

app-service

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

SEE ALSO

glob, list, regex, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

security log profile

NAME

profile - Configures a Security log profile.

MODULE

security log

SYNTAX

Configure the profile component within the security log module using the syntax shown in the following sections.

CREATE/MODIFY

```
create profile [name]
modify profile [name]
options:
  antifraud [none | add | delete | modify | replace-all-with] {
    name [string] {
  encode-fields [none | add | delete | replace-all-with] { [integer] ... }
  events [none | add | delete | modify | replace-all-with] {
    type [alert | login] {
```

```

format {
  type [none | default | user-defined]
  user-template [string]
}
rate-limit [integer]
}
rate-limit-template [string]
remote-publisher [[name] | none]
}
}
app-service [[string] | none]
application [none | add | delete | modify | replace-all-with] {
  name [string] {
options:
facility [local0 | local1 | local2 | local3 | local4 | local5 | local6 | local7]
filter [none | add | delete | modify | replace-all-with] {
  key [request-type | protocol | response-code | http-method |
search-all | search-in-headers | search-in-post-data | search-in-query-string | search-in-request | search-in-uri] {
  options:
values [none | add | delete | replace-all-with] { [string] ... }
}
}
format {
  field-delimiter [string]
  field-format [string]
  fields [none | { [string] ... }]
  type [predefined | user-defined]
  user-string [string]
}
guarantee-logging [enabled | disabled]
guarantee-response-logging [enabled | disabled]
local-storage [enabled | disabled]
logic-operation [and | or]
maximum-entry-length [1k | 2k | 10k | 64k]
maximum-header-size [integer]
maximum-query-size [integer]
maximum-request-size [integer]
protocol [udp | tcp | tcp-rfc3195]
remote-storage [none | remote | splunk | arcsight]
report-anomalies [enabled | disabled]
response-logging [none | illegal | all]
servers [none | add | delete | modify | replace-all-with] {
  [IPv4:port | IPv6.port ... ]
}
}
}
built-in [enabled | disabled]
description [string]
dos-application [none | add | delete | modify | replace-all-with] {
  name [string] {
options:
local-publisher [name]
remote-publisher [name]
}
}
bot-defense [none | add | delete | modify | replace-all-with] {
  name [string] {
options:
local-publisher [name]
remote-publisher [name]
filter {
  log-illegal-requests [disabled | enabled]
  log-challenged-requests [disabled | enabled]
  log-legal-requests [disabled | enabled]
  log-captcha-challenged-requests [disabled | enabled]
  log-bot-signature-matched-requests [disabled | enabled]
}
}
}
flowspec {
log-publisher [none | [name]]
}
ip-intelligence {
aggregate-rate [integer]
log-publisher [none | [name]]
log-translation-fields [disabled | enabled]
log-shun [disabled | enabled]
log-geo [disabled | enabled]
log-rtbh [disabled | enabled]
log-scrubber [disabled | enabled]
}
port-misuse {
log-publisher [none | [name]]
aggregate-rate [integer]
}
}
traffic-statistics {

```

```

log-ctive-flows [disabled | enabled]
log-publisher [none | [name]]
log-missed-flows [disabled | enabled]
log-reaped-flows [disabled | enabled]
log-syncookies [disabled | enabled]
log-syncookies-whitelist [disabled | enabled]
}
network [add | delete | modify | none | replace-all-with] {
  name [string] {
options:
filter {
  log-acl-match-accept [disabled | enabled]
  log-acl-match-drop [disabled | enabled]
  log-acl-match-reject [disabled | enabled]
  log-ip-errors [disabled | enabled]
  log-tcp-errors [disabled | enabled]
  log-tcp-events [disabled | enabled]
  log-translation-fields [disabled | enabled]
  log-geo-always [disabled | enabled]
  log-uuid-field [disabled | enabled]
}
rate-limit {
  acl-match-accept [integer]
  acl-match-drop [integer]
  acl-match-reject [integer]
  ip-errors [integer]
  tcp-errors [integer]
  tcp-events [integer]
  aggregate-rate [integer]
}
format {
  field-list [none | { acl_policy_name | acl_policy_type | acl_rule_name | acl_rule_uuid | action | bigip_hostname | context_name | context_type |
  dest_ip | dest_port | drop_reason | management_ip_address | protocol | route_domain |
  sa_translation_pool | sa_translation_type | src_ip | src_port | translated_dest_ip |
  translated_dest_port | translated_ip_protocol | translated_route_domain |
  translated_src_ip | translated_src_port | translated_vlan | vlan }]
  field-list-delimiter [string]
  type [field-list | none | user-defined]
  user-defined [string]
}
publisher [none | [name]]
}
}
nat {
  end-inbound-session [backup-allocation-only | disabled | enabled]
  errors [disabled | enabled]
  format {
end-inbound-session {
  field-list [none | { context_name | duration | route_domain | sub_id | translated_dest_port | translated_src_port | dest_ip | event_name | src_ip |
  timestamp | translated_route_domain | dest_port | protocol | src_port | translated_dest_ip | translated_src_ip}]
  field-list-delimiter [string]
  type [field-list | none | user-defined]
  user-defined [string]
}
end-outbound-session {
  field-list [none | { context_name | duration | route_domain | sub_id | translated_dest_port | translated_src_port | dest_ip | event_name | src_ip |
  timestamp | translated_route_domain | dest_port | protocol | src_port | translated_dest_ip | translated_src_ip}]
  field-list-delimiter [string]
  type [field-list | none | user-defined]
  user-defined [string]
}
}
errors {
  field-list [none | { context_name | duration | route_domain | sub_id | translated_dest_port | translated_src_port | dest_ip | event_name | src_ip |
  timestamp | translated_route_domain | dest_port | protocol | src_port | translated_dest_ip | translated_src_ip}]
  field-list-delimiter [string]
  type [field-list | none | user-defined]
  user-defined [string]
}
}
quota-exceeded {
  field-list [none | { context_name | duration | route_domain | sub_id | translated_dest_port | translated_src_port | dest_ip | event_name | src_ip |
  timestamp | translated_route_domain | dest_port | protocol | src_port | translated_dest_ip | translated_src_ip}]
  field-list-delimiter [string]
  type [field-list | none | user-defined]
  user-defined [string]
}
}
start-inbound-session {
  field-list [none | { context_name | duration | route_domain | sub_id | translated_dest_port | translated_src_port | dest_ip | event_name | src_ip |
  timestamp | translated_route_domain | dest_port | protocol | src_port | translated_dest_ip | translated_src_ip}]
  field-list-delimiter [string]
  type [field-list | none | user-defined]
  user-defined [string]
}
}
start-outbound-session {
  field-list [none | { context_name | duration | route_domain | sub_id | translated_dest_port | translated_src_port | dest_ip | event_name | src_ip |
  timestamp | translated_route_domain | dest_port | protocol | src_port | translated_dest_ip | translated_src_ip}]
  field-list-delimiter [string]
  type [field-list | none | user-defined]
}
}

```

```

user-defined [string]
}
}
log-publisher [none | [name]]
log-subscriber-id [disabled | enabled]
lsn-legacy-mode [disabled | enabled]
quota-exceeded [disabled | enabled]
rate-limit {
aggregate-rate [integer]
end-inbound-session [integer]
end-outbound-session [integer]
errors [integer]
quota-exceeded [integer]
start-inbound-session [integer]
start-outbound-session [integer]
}
start-inbound-session [backup-allocation-only | disabled | enabled]
end-outbound-session {
action [backup-allocation-only | disabled | enabled]
elements [add | delete | none | replace-all-with] destination
}
start-outbound-session {
action [backup-allocation-only | disabled | enabled]
elements [add | delete | none | replace-all-with] destination
}
}
protocol-dns [add | delete | modify | none | replace-all-with] {
name [string] {
options:
filter {
log-dns-drop [disabled | enabled]
log-dns-filtered-drop [disabled | enabled]
log-dns-malformed [disabled | enabled]
log-dns-malicious [disabled | enabled]
log-dns-reject [disabled | enabled]
}
format {
field-list [none | { action | attack_type | context_name | date_time | dest_ip | dest_port |
dns_query_name | dns_query_type | src_ip | src_port | vlan | route_domain }]
field-list-delimiter [string]
type [field-list | none | user-defined]
user-defined [string]
}
publisher [none | [name]]
}
}
protocol-dns-dos-publisher [none | [name]]
protocol-sip [add | delete | modify | none | replace-all-with] {
name [string] {
options:
filter {
log-sip-drop [disabled | enabled]
log-sip-global-failures [disabled | enabled]
log-sip-malformed [disabled | enabled]
log-sip-redirection-responses [disabled | enabled]
log-sip-request-failures [disabled | enabled]
log-sip-server-errors [disabled | enabled]
}
format {
field-list [none | { action | attack_type | context_name | date_time | dest_ip | dest_port |
sip_method_type | sip_caller | sip_callee | src_ip | src_port | vlan | route_domain }]
field-list-delimiter [string]
type [field-list | none | user-defined]
user-defined [string]
}
publisher [none | [name]]
}
}
protocol-sip-dos-publisher [none | [name]]
dos-network-publisher [none | [name]]
protocol-transfer [none | add | delete | modify | replace-all-with] {
name [string] {
options:
publisher [name]
}
}
}
edit profile [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties

DISPLAY
list profile
list profile [ [ [name] | [glob] | [regex] ] ... ]
show running-config profile
show running-config profile [ [ [name] | [glob] | [regex] ] ... ]

```

options:

- all-properties
- non-default-properties
- one-line
- partition
- recursive

DELETE

delete profile [name]

DESCRIPTION

You can use the profile component to create, modify, display, or delete a Security log profile for use with Security Logging functionality.

EXAMPLES

create profile my_log_profile

Creates a custom Security log profile named my_log_profile with initial settings.

list profile

Displays the properties of all Security log profiles.

OPTIONS

antifraud

Adds, deletes, or replaces a single Anti-Fraud Security sub-profile. You can configure the following options for Anti-Fraud Security:

encode-fields

Adds, deletes, or replaces a set of antifraud-storage-field IDs for which the system performs URL-encoding before logging.

events

Adds, deletes, or replaces a set of events (alert, login) used by the system to log data. You can configure the following options for each event:

format

Specifies a storage format in Anti-Fraud Security. You can configure the following options for the storage format:

type Specifies a type of the storage format. The options are:

default

Specifies that the log displays a predefined format and antifraud-storage-field fields.

user-defined

Specifies that the log displays any free text that you type in the user-template which can include relevant antifraud-storage-field fields for this event.

rate-limit

This option is used to set the rate for the Anti-Fraud log event that can be logged per second, per virtual-server (per TMM).

user-template

Specifies a user template in the user-defined storage format.

rate-limit-template

Specifies a template for rate-limit event logging.

remote-publisher

Specifies the name of the log publisher used for logging Anti-Fraud events.

app-service

Specifies the name of the application service to which the profile belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the profile. Only the application service can modify or delete the profile.

application

Adds, deletes, or replaces a single Application Security sub-profile. You can configure the following options for Application Security:

facility

Specifies the facility category of the logged traffic in Application Security. Select between local0 and local7.

filter

Adds, deletes, or replaces a set of request filters in Application Security. You can configure the following options for a request filter:

key Specifies a unique key for the request filter. This option is required for the operations create, delete, modify, and replace-all-with. The options are:

request-type

Specifies which kind of requests the system, or server, logs.

protocol

Specifies whether request logging is dependent on the protocol.

`response-code`
Specifies whether request logging is dependent on the response status code.

`http-method`
Specifies whether request logging is dependent on the HTTP method.

`search-all`, `search-in-headers`, `search-in-post-data`, `search-in-query-string`, `search-in-request`,
`search-in-uri`
Specifies whether the request logging is dependent on a specific string, and if so, the part of the request where the system must find the string. You can select only one of these filters, the default is `search-all`, which means that the system logs all requests, regardless of string.

`values`
Adds, deletes, or replaces a set of values in the request filter.

`format`
Specifies a storage format in Application Security. You can configure the following options for the storage format:

`field-delimiter`
Specifies a field delimiter in the predefined storage format. You may not use the % character. The default delimiter is the comma character, for CSV.

`field-format`
Specifies a field format (for each key/value pair) in the predefined storage format. Use %k for key and %v for value. The default format is empty that is interpreted as "%v", for CSV.

`fields`
Replaces a set of fields in the predefined storage format. The order in the set is important - the server displays the selected traffic items in the log sequentially according to it.

`type` Specifies a type of the storage format. The options are:

`predefined`
Specifies that the log displays only the predefined items you select in the fields.

`user-defined`
Specifies that the log displays any free text that you type in the user-string which can include the predefined items.

`user-string`
Specifies a user string in the user-defined storage format.

`guarantee-logging`
Indicates whether to guarantee local logging in Application Security.

`guarantee-response-logging`
Indicates whether to guarantee local response logging in Application Security. In order to enable it, you must first enable `guarantee-logging`, and set `response-logging` to either `illegal` or `all`.

`local-storage`
Enables or disables local storage in Application Security.

`logic-operation`
Specifies the logic operation on the associated filters in Application Security. The options are:

`and` Specifies that requests must pass all filters in order for the system, or server, to log the requests.

`or` Specifies that requests must meet at least one filter in order for the system, or server, to log the requests. This is the default value.

`maximum-entry-length`
Specifies the maximum entry length in Application Security. The options are:

`1k` This is the possible length for remote servers that support the `udp` protocol.

`2k` This is the default length for remote servers that support the `tcp`, `udp` and `tcp-rfc3195` protocols.

`10k`, `64k`
These are possible lengths for remote servers that support the `tcp` and `udp` protocol.

`maximum-header-size`
Specifies the maximum headers size in Application Security.

`maximum-query-size`
Specifies the maximum query string size in Application Security.

`maximum-request-size`
Specifies the maximum request size in Application Security.

`name` Specifies a dummy name for enabled Application Security. This option is required for the operations `create`, `delete`, `modify`, and `replace-all-with`.

`protocol`
Specifies the protocol supported by the remote server in Application Security. Select either: `tcp` (the default value), `udp`, or `tcp-rfc3195`.

`remote-storage`
Specifies a remote storage type in Application Security. The options are:

`none` Specifies that the system does not store traffic on any remote logging server.

`remote`
Specifies that the system stores all traffic on a remote logging server, like a `syslog`.

`splunk`
Specifies that the system stores all traffic on a reporting server (Splunk) using a preconfigured storage format. Key/value pairs are used in the log messages.

`arcsight`
Specifies that the system stores all traffic on a remote logging server using the predefined ArcSight settings for the logs. The log messages are in Common Event Format (CEF).

`report-anomalies`
Indicates whether to report detected anomalies in Application Security.

`response-logging`
Specifies a response logging type in Application Security. The options are:

`none` Specifies that the system does not log responses. This is the default value.

`illegal`
Specifies that the system logs responses to illegal requests.

`all` Specifies that the system logs all responses if the associated request-type filter has the `all` value.

`servers`
Adds, deletes, or replaces a set of remote servers in Application Security, by specifying an IP address and service port in the format `[IPv4:port]` or `[IPv6.port]`.

`built-in`
Displays whether this profile is predefined or user-defined.

`description`
User defined description.

`dos-application`
Adds, deletes, or replaces a single DoS (Application) Protection sub-profile. You can configure the following options for DoS (Application) Protection:

`local-publisher`
Specifies the name of the local log publisher used for Application DoS attacks. Note: This publisher should have a single `local-database` destination.

`name` Specifies a dummy name for enabled DoS (Application) Protection. This option is required for the operations `create`, `delete`, `modify`, and `replace-all-with`.

`remote-publisher`
Specifies the name of the remote log publisher used for Application DoS attacks. Note: This publisher should have `arcsight` or `splunk` destinations.

`bot-defense`
Adds, deletes, or replaces a single Bot Defense sub-profile. You can configure the following options for Bot Defense:

`name` Specifies a dummy name for enabled Bot Defense. This option is required for the operations `create`, `delete`, `modify`, and `replace-all-with`.

`local-publisher`
Specifies the name of the local log publisher used for Bot Defense log messages. Note: This publisher should have a single `local-database` destination.

`remote-publisher`
Specifies the name of the remote log publisher used for Bot Defense log messages. Note: This publisher should have only `splunk` destinations.

`filter`
Following options are available which enable or disable the logging of Bot Defense log messages:

`log-illegal-requests`
This option is used to enable or disable the logging of illegal requests.

`log-challenged-requests`
This option is used to enable or disable the logging of challenged requests.

`log-legal-requests`
This option is used to enable or disable the logging of legal requests.

log-captcha-requests

This option is used to enable or disable the logging of captcha challenged requests.

log-bot-signature-matched-requests

This option is used to enable or disable the logging of reported bot signature requests. =back

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

flowspec

Security FlowSpec log configuration

log-publisher

Specifies the name of the log publisher used for Security FlowSpec log events.

ip-intelligence

You can configure the following options under this:

aggregate-rate

This option is used to set the aggregate rate limit that applies to any ip intelligence log message.

log-publisher

Specifies the name of the log publisher used for IP Intelligence events.

log-translation-fields

This option is used to enable or disable the logging of translated (i.e server side) fields in IP Intelligence log messages. Translated fields include (but not limited to) Source Address/Port, Destination Address/Port, IP Protocol, Route Domain and Vlan.

log-shun

This option is used to enable or disable the logging of shun IP Intelligence events.

log-geo

This option is used to enable or disable the logging of geo location in shun IP Intelligence event.

log-rtbh

This option is used to enable or disable the logging of rtbh IP Intelligence events.

log-scrubber

This option is used to enable or disable the logging of scrubber IP Intelligence events.

port-misuse

You can configure the following options under this:

log-publisher

Specifies the name of the log publisher used for port misuse events.

aggregate-rate

This option is used to set the rate limit that applies to any port misuse log messages.

traffic-statistics

You can configure the following options under this:

log-active-flows

This option is used to enable and disable the logging of number of active flows on client side. The number of flows are logged globally, per virtual server and per route domain periodically if number of active flows increased or decreased.

log-publisher

Specifies the name of the log publisher used for Traffic Statistics logs.

log-reaped-flows

This option is used to enable and disable the logging of number of reaped flows on client side. The number of flows are logged globally, per virtual server and per route domain periodically if number of active flows increased or decreased.

log-missed-flows

This option is used to enable and disable the logging of number of TCP packets (non SYN/ACK) were dropped because of the flow table lookup failed. The number of packets are logged globally, and per route domain periodically.

log-syncookies

This option is used to enable and disable the logging of number of syncookies generated, accepted and rejected in the context globally and per virtual server. These log messages will be generated periodically.

log-syncookies-whitelist

This option is used to enable and disable the logging of number of syncookies whitelist hits, accepted and rejected in the context globally and per virtual server. These log messages will be generated periodically.

network

Add, delete, modify or replace a single Network Security sub-profile. You can configure the following options under this:

filter

Following options are available which enable or disable the logging of corresponding Network events:

log-acl-match-accept

This option is used to enable or disable the logging of packets that match ACL rules configured with action = Accept or action = Accept Decisively.

log-acl-match-drop

This option is used to enable or disable the logging of packets that match ACL rules configured with action = Drop.

log-acl-match-reject

This option is used to enable or disable the logging of packets that match ACL rules configured with action = Reject.

log-ip-errors

This option is used to enable or disable the logging of IP error packets.

log-tcp-errors

This option is used to enable or disable the logging of TCP error packets.

log-tcp-events

This option is used to enable or disable the logging of TCP events on client side. Only 'Established' and 'Closed' states of a TCP session are logged if this option is enabled.

log-translation-fields

This option is used to enable or disable the logging of translated (i.e server side) fields in ACL match and TCP events. Translated fields include (but not limited to) Source Address/Port, Destination Address/Port, IP Protocol, Route Domain and Vlan.

log-geo-always

This option is used to enable or disable the logging of Geographic IP Location information fields in ACL match and TCP logging. Geographic information includes the country code of Source Address and Destination Address.

log-uuid-field

This option is used to enable or disable the logging of ACL rule UUID field in ACL match and TCP logging. If the `acl_rule_uuid` field is explicitly specified in field-list or user-defined formats, UUID value will be logged regardless of state of this option.

rate-limit

Following options are available to set throttling rate limits for the corresponding logging network events:

acl-match-accept

This option is used to set rate limits for the logging of packets that match ACL rules configured with action = Accept or action = Accept Decisively. This option is effective only if logging of this message type is enabled.

acl-match-drop

This option is used to set rate limits for the logging of packets that match ACL rules configured with action = Drop. This option is effective only if logging of this message type is enabled.

acl-match-reject

This option is used to set rate limits for the logging of packets that match ACL rules configured with action = Reject. This option is effective only if logging of this message type is enabled.

ip-errors

This option is used to set rate limits for the logging of IP error packets. This option is effective only if logging of this message type is enabled.

tcp-errors

This option is used to set rate limits for the logging of TCP error packets. This option is effective only if logging of this message type is enabled.

tcp-events

This option is used to set rate limits for the logging of TCP events on client side. This option is effective only if logging of this message type is enabled.

aggregate-rate

This option is used to set the aggregate rate limit that applies to any network logging message.

format

Specifies the Storage format in Network Security sub-profile. These settings are only used to format the log messages destined to a Remote Syslog server. You can configure the following options for the storage format:

field-list

Specifies a set of fields to be logged. This option is valid when storage format type is field-list. The order in the set is important - the server displays the selected traffic items in the log sequentially according to it. User can pick fields from the following list:
`acl_policy_name, acl_policy_type, acl_rule_name, acl_rule_uuid, action, bigip_hostname, context_name, context_type, date_time, dest_fqdn, dest_geo, dest_ip, dest_port, drop_reason, management_ip_address, protocol, route_domain, sa_translation_pool, sa_translation_type, source_fqdn, source_user, src_geo, src_ip, src_port, translated_dest_ip, translated_dest_port, translated_ip_protocol, translated_route_domain, translated_src_ip, translated_src_port,`

translated_vlan, vlan.

field-list-delimiter

Specifies the delimiter string in field-list storage format type. The default delimiter is the comma character, for CSV. This option is valid when storage format type is field-list. Special character \$ should not be used in delimiter string as it is reserved for internal usage. Also, the maximum length allowed for field-list-delimiter is 31 characters (excluding NUL terminator).

type Specifies a type of the storage format. The options are:

field-list

Specifies that the log displays only the items you specify in the field-list with field-list-delimiter as the delimiter between the items.

none Default format type. With this option, the messages will be logged in the following format:

```
"management_ip_address","bigip_hostname","context_type","context_name","src_geo","src_ip", "dest_geo","dest_ip","src_port","dest_por
```

user-defined

Specifies that the log displays the message as per the user-defined string format.

user-defined

Specifies the format of log message in form of user defined string. This option is valid when storage format type is user-defined. Maximum configurable length is 512 characters. Any of the following items, if wrapped within \${ }, will be substituted with the actual value when generating the log: acl_policy_name, acl_policy_type, acl_rule_name, acl_rule_uuid, action, bigip_hostname, context_name, context_type, date_time, dest_fqdn, dest_geo, dest_ip, dest_port, drop_reason, management_ip_address, protocol, route_domain, sa_translation_pool, sa_translation_type, source_fqdn, source_user, src_geo, src_ip, src_port, translated_dest_ip, translated_dest_port, translated_ip_protocol, translated_route_domain, translated_src_ip, translated_src_port, translated_vlan, vlan.

publisher

Specifies the name of the log publisher used for Network events.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

nat This section is used to configure log settings related to events applicable to firewall NAT feature. Following options are available under this section:

end-inbound-session

Event for end of incoming connection to a translated address. Inbound connections are supported only for dynamic-pat source translation. Following options can be configured for logging this event:

backup-allocation-only

Enable logging this event when translation is done using backup address in the source translation object configured in dynamic-pat mode. This is only applicable when lsn-legacy-mode is enabled.

disabled

Disables logging this event.

enabled

Enables logging this event when translation is done using primary address or backup address in the source translation object.

errors

Event for errors encountered while attempting source or destination translation.

disabled

Disables logging for this event.

enabled

Enables logging for this event.

log-publisher

Specifies the name of log publisher used to log NAT related events to one (or more) remote or local destinations.

lsn-legacy-mode

Specifies whether translation events (and other NAT events) are logged in existing CGNAT/LSN formats (for backward compatibility with LSN events).

log-subscriber-id

When enabled, the subscriber ID associated with a subscriber IP address will be printed in the logs.

quota-exceeded
Event for when client exceeded allocated resource limit.

disabled

Disables logging for this event.

enabled

Enables logging for this event.

rate-limit
Following options are available to set throttling rate limits for the corresponding logging FW NAT events:

aggregate-rate-limit
This option is used to set the aggregate rate for all the FW NAT log events that can be logged per second.

end-inbound-session
This option is used to rate limit the end inbound session log events per second.

end-outbound-session
This option is used to rate limit the end outbound session log events per second.

errors
This option is used to rate limit the errors to be logged per second.

start-inbound-session
This option is used to rate limit the start inbound session log events per second.

start-outbound-session
This option is used to rate limit the start outbound session log events per second.

quota-exceeded
This option is used to rate limit the quota exceeded log events per second.

start-inbound-session
Event for start of incoming connection to a translated address. Inbound connections are supported only for dynamic-pat source translation. Following options can be configured for logging this event:

backup-allocation-only

Enable logging this event when translation is done using backup address in the source translation object configured in dynamic-pat mode.

disabled

Disables logging this event.

enabled

Enables logging this event when translation is done using primary address or backup address in the source translation object.

end-outbound-session
Event for end of outbound translation session, when outbound flow is deleted.

action
Specifies what action is taken at the time of logging the event. Possible options are: backup-allocation-only, disabled and enabled.

elements
Optional elements that can be logged for the event. This is applicable only if lsn-legacy-mode is enabled.

destination

Optional element, if selected, is used to log destination address and port in the applicable log event.

start-outbound-session
Event for start of outbound translation session, when outbound flow is created.

action
Specifies what action is taken at the time of logging the event. Possible options are: backup-allocation-only, disabled and enabled.

elements
Optional elements that can be logged for the event. This is applicable only if lsn-legacy-mode is enabled.

destination

Optional element, if selected, is used to log destination address and port in the applicable log event.

protocol-dns

Add, delete, modify or replace a single Protocol (DNS) Security sub-profile. You can configure the following options under this:

filter

Following options are available which enable or disable the logging of corresponding Network events:

log-dns-drop

This option is used to enable or disable the logging of dropped DNS packets.

log-dns-filtered-drop

This option is used to enable or disable the logging of DNS packets that are dropped due to filtering.

log-dns-malformed

This option is used to enable or disable the logging of malformed DNS packets.

log-dns-malicious

This option is used to enable or disable the logging of malicious DNS packets.

log-dns-reject

This option is used to enable or disable the logging of rejected DNS packets.

format

Specifies the Storage format in Protocol (DNS) Security sub-profile. These settings are only used to format the log messages destined to a Remote Syslog server. You can configure the following options for the storage format:

field-list

Specifies a set of fields to be logged. This option is valid when storage format type is field-list. The order in the set is important - the server displays the selected traffic items in the log sequentially according to it. User can pick fields from the following list: action, attack_type, context_name, date_time, dest_ip, dest_port, dns_query_name, dns_query_type, src_ip, src_port, vlan.

field-list-delimiter

Specifies the delimiter string in field-list storage format type. The default delimiter is the comma character, for CSV. This option is valid when storage format type is field-list. Special character \$ should not be used in delimiter string as it is reserved for internal usage. Also, the maximum length allowed for field-list-delimiter is 31 characters (excluding NUL terminator).

type Specifies a type of the storage format. The options are:

field-list

Specifies that the log displays only the items you specify in the field-list with field-list-delimiter as the delimiter between the items.

none Default format type. With this option, the messages will be logged in the following format:

```
"date_time", "context_name", "vlan", "dns_query_type", "dns_query_name", "attack_type",  
"action", "src_ip", "dest_ip", "src_port", "dest_port", "route_domain"
```

user-defined

Specifies that the log displays the message as per the user-defined string format.

user-defined

Specifies the format of log message in form of user defined string. This option is valid when storage format type is user-defined. Maximum configurable length is 512 characters. Any of the following items, if wrapped within \${ }, will be substituted with the actual value when generating the log: action, attack_type, context_name, date_time, dest_ip, dest_port, dns_query_name, dns_query_type, route_domain, src_ip, src_port, vlan.

name Specifies a dummy name for enabled Protocol (DNS) Security. This option is required for the operations create, delete, modify, and replace-all-with.

publisher

Specifies the name of the log publisher used for DNS events.

protocol-dns-dos-publisher

Specifies the name of the log publisher used for DNS DoS events.

dos-network-publisher

Specifies the name of the log publisher used for DoS Network events.

protocol-sip

Add, delete, modify or replace a single Protocol (SIP) Security sub-profile. You can configure the following options under this:

filter

Following options are available which enable or disable the logging of corresponding protocol sip events:

log-sip-drop

This option is used to enable or disable the logging of dropped SIP packets.

log-sip-global-failures

This option is used to enable or disable the logging of SIP packets that resulted in global failures.

log-sip-malformed

This option is used to enable or disable the logging of malformed SIP packets.

log-sip-redirection-responses

This option is used to enable or disable the logging of SIP packets that resulted in sending redirection response.

log-sip-request-failures

This option is used to enable or disable the logging of SIP request failures.

log-sip-server-errors

This option is used to enable or disable the logging of SIP packets that resulted in server errors.

format

Specifies the Storage format in Protocol (SIP) Security sub-profile. These settings are only used to format the log messages destined to a Remote Syslog server. You can configure the following options for the storage format:

field-list

Specifies a set of fields to be logged. This option is valid when storage format type is field-list. The order in the set is important - the server displays the selected traffic items in the log sequentially according to it. User can pick fields from the following list: action, attack_type, context_name, date_time, dest_ip, dest_port, dns_query_name, dns_query_type, src_ip, src_port, vlan.

field-list-delimiter

Specifies the delimiter string in field-list storage format type. The default delimiter is the comma character, for CSV. This option is valid when storage format type is field-list. Special character \$ should not be used in delimiter string as it is reserved for internal usage. Also, the maximum length allowed for field-list-delimiter is 31 characters (excluding NUL terminator).

type Specifies a type of the storage format. The options are:

field-list

Specifies that the log displays only the items you specify in the field-list with field-list-delimiter as the delimiter between the items.

none Default format type. With this option, the messages will be logged in the following format:

```
"date_time", "context_name", "vlan", "sip_method_type", "sip_caller", "sip_callee",  
"attack_type", "action", "src_ip", "dest_ip", "src_port", "dest_port", "route_domain"
```

user-defined

Specifies that the log displays the message as per the user-defined string format.

user-defined

Specifies the format of log message in form of user defined string. This option is valid when storage format type is user-defined. Maximum configurable length is 512 characters. Any of the following items, if wrapped within \${ }, will be substituted with the actual value when generating the log: action, attack_type, context_name, date_time, dest_ip, dest_port, dns_query_name, dns_query_type, route_domain, src_ip, src_port, vlan.

name Specifies a dummy name for enabled Protocol (SIP) Security. This option is required for the operations create, delete, modify, and replace-all-with.

publisher

Specifies the name of the log publisher used for SIP events.

protocol-sip-dos-publisher

Specifies the name of the log publisher used for SIP DoS events.

protocol-transfer

Adds, deletes, or replaces a single Protocol (Transfer) Security sub-profile. You can configure the following options for Protocol (Transfer) Security:

name Specifies a dummy name for enabled Protocol (Transfer) Security. This option is required for the operations create, delete, modify, and replace-all-with.

publisher

Specifies the name of the log publisher used for Protocol Security log messages. Note: This publisher should have either local-database, local-syslog, remote-syslog, arcsight or splunk single destination.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

asm http-method, asm response-code, create, delete, edit, glob, list, ltm virtual, modify, regex, security, security log, security log storage-field, show, sys log-config destination, sys log-config publisher, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2018-11-27 security log profile(1)

security log protocol-dns-storage-field

NAME

protocol-dns-storage-field - Lists the available storage format fields that can be used in the context of Protocol DNS Security Logging.

MODULE

security log

SYNTAX

Retrieve the list of the protocol-dns-storage-field values using the syntax shown in the following sections.

DISPLAY

```
list protocol-dns-storage-field
list protocol-dns-storage-field [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  one-line
  app-service
```

DESCRIPTION

Use this command to display the possible values of the protocol-dns-storage-field object to be used in the context of Protocol DNS Security Logging. These possible values are predefined traffic items available for the server to log in the context of DNS event logging (for example, Malformed, Malicious, or Dropped DNS packets).

EXAMPLES

```
list protocol-dns-storage-field
```

Displays all the storage fields supported by Protocol DNS Security Logging.

OPTIONS

app-service
Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

SEE ALSO

glob, list, regex, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-12-20 security log protocol-dns-storage-field(1)

security log protocol-sip-storage-field

NAME

protocol-sip-storage-field - Lists the available storage format fields that can be used in the context of Protocol SIP Security Logging.

MODULE

security log

SYNTAX

Retrieve the list of the protocol-sip-storage-field values using the syntax shown in the following sections.

DISPLAY

```
list protocol-sip-storage-field
list protocol-sip-storage-field [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  one-line
  app-service
```

DESCRIPTION

Use this command to display the possible values of the protocol-sip-storage-field object to be used in the context of Protocol SIP Security Logging. These possible values are predefined traffic items available for the server to log in the context of SIP event logging (e.g Dropped SIP packets).

EXAMPLES

```
list protocol-sip-storage-field
```

Displays all the storage fields supported by Protocol SIP Security Logging.

OPTIONS

```
app-service
```

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

SEE ALSO

glob, list, regex, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-03-21 security log protocol-sip-storage-field(1)

security log remote-format

NAME

remote-format - Lists the log format for different remote destinations (such as ArcSight, Splunk etc.) used by various Firewall events (such as Network, IP Intelligence, DoS etc.).

MODULE

security log

SYNTAX

Retrieve the list of the remote-format using the syntax shown in the following sections.

DISPLAY

```
list remote-format
list remote-format [ [ [name] | [glob] | [regex] ] ... ]
options:
  all
  all-properties
  app-service
  format
  one-line
```

DESCRIPTION

Use this command to display the actual log format used to send firewall event logs to remote destinations such as ArcSight, Splunk and Syslog. These log formats are used by the log destinations of log publisher configured in different sub-profiles (for example Network, IP Intelligence, DNS, DNS DoS etc.) of a security log profile.

EXAMPLES

```
list remote-format
```

Displays the log format for all firewall events.

```
list remote-format network-arcsight
```

Displays the format for Network log events (such as ACL matches, TCP events etc.) sent to an ArcSight destination.

```
list remote-format network-dos-splunk
```

Displays the format for Network DoS log events sent to a Splunk destination.

list remote-format ip-intelligence-syslog-default

Displays the format for IP Intelligence log events sent to a remote syslog destination.

OPTIONS

app-service

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

format

Displays the remote log format used by the object.

SEE ALSO

glob, list, regex, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-02-20 security log remote-format(1)

security log storage-field

NAME

storage-field - Lists the available storage format fields that can be used in the context of Application Security Logging.

MODULE

security log

SYNTAX

Retrieve the list of the storage-field values using the syntax shown in the following sections.

DISPLAY

list storage-field

list storage-field [[[name] | [glob] | [regex]] ...]

options:

all
app-service
format
id
one-line

DESCRIPTION

Use this command to display the possible values of the storage-field object to be used in the context of Application Security Logging. These possible values are predefined traffic items available for the server to log. The traffic items appear in the final format string as arguments in the printf() function, i.e. "%\$", therefore each storage field has its fixed format (specifier) and id (position).

EXAMPLES

list storage-field

Displays all the storage fields supported by Application Security Logging.

OPTIONS

app-service

Displays the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

format

Displays a format of the field (s - string, d - decimal). It corresponds to the conversion specifier in the printf() function.

id Displays an order ID of the field (starting from 1). It corresponds to the position in the argument list of the desired argument in the printf() function.

SEE ALSO

glob, list, regex, security log profile, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

security malicious-sources device-ids

NAME

device-ids - Displays the malicious device IDs.

MODULE

security malicious-sources

SYNTAX

Display information about the device-ids component within the security malicious-sources module using the following syntax.

DISPLAY

show device-ids

options:

field-fmt

DESCRIPTION

You can use the device-ids component to display the malicious device IDs.

EXAMPLES

show device-ids

Display the list of detected malicious device IDs.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

security malicious-sources ip-addresses

NAME

ip-addresses - Displays the malicious source IP addresses.

MODULE

security malicious-sources

SYNTAX

Display information about the ip-addresses component within the security malicious-sources module using the following syntax.

DISPLAY

show ip-addresses

options:

field-fmt

DESCRIPTION

You can use the ip-addresses component to display the malicious source IP addresses.

EXAMPLES

show ip-addresses

Display the list of detected malicious source IP addresses.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO
show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016-2017. All rights reserved.

BIG-IP 2017-06-22 security malicious-sources ip-addresses(1)

security nat destination-translation

NAME

destination-translation - Configures a Security NAT destination translation object.

MODULE

security nat

SYNTAX

CREATE/MODIFY

create destination-translation [name]

modify destination-translation [name | all]

options:

addresses [add | delete | modify | none | replace-all-with] {

[[ip address] [ip prefix] [ip range]]

}

app-service [[string] | none]

description [string]

ports [add | delete | modify | none | replace-all-with] {

[[port] [port-range]]

}

type [static-nat | static-pat]

edit destination-translation [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list destination-translation

list destination-translation [[[name] | [glob] | [regex]] ...]

show running-config destination-translation

show running-config destination-translation [[[name] | [glob] | [regex]] ...]

options:

all

all-properties

non-default-properties

one-line

recursive

show destination-translation

show destination-translation [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

detail

field-fmt

DELETE

delete destination-translation [name | all]

DESCRIPTION

A destination-translation NAT object is a set of IP Address(es) and port numbers that the BIG-IP system uses as public-side addresses and ports. When this object is assigned to a Security NAT Policy rule (which is associated to a virtual server), any incoming traffic to this virtual server that matches the rule, will have their private destination addresses (and/or ports) translated to a public address and/or port from this destination-translation object.

EXAMPLES

```
create destination-translation d1 type static-pat addresses add { 10.10.10.0/24 } ports add { 7000-8000 }
```

Creates the destination-translation object d1 that contains the translation addresses in the range of 10.10.10.0/24, translation port range 7000-8000 and uses static-pat as the translation algorithm.

```
delete destination-translation d1
```

Deletes the destination-translation named d1.

OPTIONS

addresses

Specifies the set of translation IP addresses available. This is a collection of either one (or more) IP Address, one (or more) IP prefixes with their prefix lengths and/or IP Address range. All public-side (destination) addresses come from the subnets you enter in this property.

app-service

Specifies the name of the application service to which this object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the destination-translation component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

ports

Specifies the range of port numbers available for use with translation IP addresses.

type Specifies which kind of translation is performed. Available options are static-nat and static-pat.

static-nat

Using this translation type in the destination-translation, only (destination) IP Address translation is performed (and no port translation) for the incoming client traffic that matches the NAT Policy Rule using this translation object. There is a static (pre-defined) 1:1 mapping between the untranslated IP Address(es) and the translated IP Address(es) specified in this object.

static-pat

Using this translation type in the destination-translation, (destination) port translation is performed for the incoming client traffic that matches the NAT Policy Rule using this translation object. There is a static (pre-defined) 1:1 mapping between the untranslated port(s) and the translated port(s) specified in this object. In addition, if translation address(es) are specified in the destination-translation, it also performs IP Address translation (in the same fashion as done for static-nat).

SEE ALSO

security nat policy, security nat source-translation, ltm virtual, create, delete, edit, glob, list, ltm, modify, regex, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-03-14 security nat destination-translation(1)

security nat policy

NAME

policy - Configures nat policy.

MODULE

security nat

SYNTAX

Modify the policy component within the security nat module using the syntax shown in the following sections.

CREATE/MODIFY

create policy [name]

modify policy [name]

options:

app-service [[string] | none]

description [string]

rules [add | delete | modify | replace-all-with] {

[[name]] {

options:

app-service [[string] | none]

```

description [string]
ip-protocol [protocol name]
log-profile [name | none]
place-after [first | last | [rule name]]
place-before [first | last | [rule name]]
status [disabled | enabled]
destination {
  address-lists [add | default | delete | replace-all-with] {
    [address list names...]
  }
  address-lists none
  addresses [add | default | delete | replace-all-with] {
    [ [ip address] | [ip address/prefixlen] ]
  }
  addresses none
  port-lists [add | default | delete | replace-all-with] {
    [port list names...]
  }
  port-lists none
  ports [add | default | delete | none | replace-all-with] {
    [ [port] | [port1-port2] ]
  }
  ports none
  proxy-arp [enabled | disabled]
  route-advertisement [enabled | disabled]
}
source {
  address-lists [add | default | delete | replace-all-with] {
    [address list names...]
  }
  address-lists none
  addresses [add | default | delete | replace-all-with] {
    [ [ip address] | [ip_address/prefixlen] ]
  }
  addresses none
  port-lists [add | default | delete | replace-all-with] {
    [port list names...]
  }
  port-lists none
  ports [add | default | delete | replace-all-with] {
    [ [port] | [port1-port2] ]
  }
  ports none
  vlans [add | default | delete | replace-all-with] {
    [vlan names...]
  }
  vlans none
}
translation {
  destination [name | none]
  source [name | none]
}
next-hop {
  gw [ip address]
  vlan [name | none]
  pool [name | none]
  type [default | pool | gateway | vlan]
}
}
rules none

```

```

edit policy
options:
  all-properties
  non-default-properties

```

```

DISPLAY
list policy
show running-config policy
options:
  all-properties
  non-default-properties
one-line

```

DESCRIPTION

You can use the policy component to configure a shareable and reusable set of nat rules which can be associated with a number of configuration objects of the following types: ltm virtual, security device-context, net route-domain.

EXAMPLES

```

create policy p1 rules add {
r1 {
  place-before first
  ip-protocol tcp
  source {
    addresses replace-all-with { 192.168.10.0/24 }
  }
}
}

```

```

    ports replace-all-with { 10000-19999 }
  }
  destination {
    addresses replace-all-with { 10.10.10.0/24 }
    ports replace-all-with { 80 443 }
  }
  translation {
    destination my_dest_nat
    source my_src_nat
  }
  next-hop {
    gw 10.10.10.10
    pool pool1
    vlan internal
    type vlan
  }
} }

```

Creates a rule entry at the beginning of the policy that matches incoming TCP traffic with source address in the range 192.168.10.0/24, source port in the range 10000-19999, destination address in the range 10.10.10.0/24, destination port 80 or 443 and if matches, performs the source translation as per source-translation object named my_src_nat and destination translation as per destination-translation object named my_dest_nat.

modify policy p1 rules delete r1

Removes the rule r1 from the policy p1.

list policy

Displays the current list of policy rules.

OPTIONS

description
User defined description.

rules
Adds, deletes, or replaces a NAT rule.

description
User defined description.

destination
address-lists
Specifies a list of address lists (see security firewall address-list) against which the packet will be compared.

addresses
Specifies a list of addresses and networks against which the packet will be compared.

port-lists
Specifies a list of port lists (see security firewall port-list) against which the packet will be compared.

ports
Specifies a list of ports and port ranges against which the packet will be compared.

proxy-arp
Enable or disable proxy arp for pre-translation destination addresses.

route-advertisement
Enable or disable route advertisements for pre-translation destination addresses.

ip-protocol
Specifies the IP protocol against which the packet will be compared.

log-profile
Specifies the name of the log profile (see security log profile) that is used to log the translation events triggered by this NAT rule.

place-after
Specifies that a new rule should be placed after another rule, first or last. If individual rules are being added (as opposed to specifying replace-all-with) then place-before or place-after must be specified.

place-before
Specifies that a new rule should be placed before another rule, first or last. If individual rules are being added (as opposed to specifying replace-all-with) then place-before or place-after must be specified.

source
address-lists
Specifies a list of address lists (see security firewall address-list) against which the packet will be compared.

addresses
Specifies a list of addresses and networks against which the packet will be compared.

port-lists

Specifies a list of port lists (see security firewall port-list) against which the packet will be compared.

ports

Specifies a list of ports and port ranges against which the packet will be compared.

vlan

Specifies a list of vlans, vlan groups and tunnels against which the packet will be compared.

next-hop

Specifies next-hop configuration for NAT rule. All these attributes are mutually exclusive.

gw Specifies a gateway address for the route.

vlan Specifies VLAN name (can be VLAN or VLAN group)

pool Specifies a gateway pool, which allows multiple, load-balanced gateways to be used for the route.

type Specifies which option to consider when multiple options were provided under next-hop.

status

Specifies whether the rule is enabled or disabled. A rule that is enabled is always checked. A rule that is disabled is never checked.

translation

Specifies the translation objects.

destination

Specifies the name of destination translation object (see security nat destination-translation). If specified, it is used to perform the destination address/port translation as per its settings. If not specified, the matching traffic's destination address/port are not translated.

source

Specifies the name of source translation object (see security nat source-translation). If specified, it is used to perform the source address/port translation as per its settings. If not specified, the matching traffic's source address/port are not translated.

SEE ALSO

create, edit, list, modify, security firewall address-list, security firewall port-list, security nat destination-translation, security nat source-translation, security log profile, tms, security device-context nat-policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015. All rights reserved.

BIG-IP 2019-08-27 security nat policy(1)

security nat source-translation

NAME

source-translation - Configures a Security NAT source translation object.

MODULE

security nat

SYNTAX

CREATE/MODIFY

```
create source-translation [name]
modify source-translation [name | all]
options:
  addresses [add | delete | modify | none | replace-all-with] {
  [ [ip address] [ip prefix] [ip range] ]
  }
  app-service [[string] | none]
  backup-addresses
  [add | delete | replace-all-with] {
  [ip address/prefix length] ...
  }
  client-connection-limit [integer value]
  description [string]
  egress-interfaces
```

```

    [add | delete | replace-all-with] {
[interface name] ...
    }
    egress-interfaces-disabled
    egress-interfaces-enabled
    exclude-addresses [add | delete | modify | none | replace-all-with] {
[ ip address] [ip prefix] [ip range] ]
    }
    exclude-address-lists [add | default | delete | none | replace-all-with] {
[address list names...]
    }
    hairpin-mode [enabled | disabled]
    icmp-echo [enabled | disabled]
    inbound-mode [endpoint-independent-filtering | explicit | none]
    eif-timeout [integer]
    pat-mode [deterministic | napt | pba]
    pcp {
        profile [ name | none ]
        selfip [ name | none]
        dslite_tunnel [ name | none ]
    }
    ports [add | delete | modify | none | replace-all-with] {
[ [port] [port-range] ]
    }
    proxy-arp [enabled | disabled]
    route-advertisement [enabled | disabled]
    type [dynamic-pat | static-nat | static-pat]
    mapping {
        mode [address-pooling-paired | endpoint-independent-mapping | none]
        timeout [integer]
    }
    port-block-allocation {
        block-idle-timeout [integer]
        block-lifetime [integer]
        block-size [integer]
        client-block-limit [integer]
        zombie-timeout [integer]
    }
}

```

edit source-translation [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list source-translation

list source-translation [[[name] | [glob] | [regex]] ...]

show running-config source-translation

show running-config source-translation [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line

show source-translation

show source-translation [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
detail
field-fmt

DELETE

delete source-translation [name | all]

DESCRIPTION

A source-translation NAT object is a set of IP Address(es) and port numbers that the BIG-IP system uses as public-side addresses and ports. When this object is assigned to a Security NAT Policy rule (which is associated to a virtual server), any incoming traffic to this virtual server that matches the rule, will have their private source addresses (and/or ports) translated to a public address and/or port from this source-translation object.

EXAMPLES

```
create source-translation s1 type dynamic-pat pat-mode napt mapping { mode endpoint-independent-mapping
timeout 600 } addresses add { 10.10.20.0/24 } ports add { 4000-5000 } client-connection-limit 100
```

Creates the source-translation object named s1 that contains the translation addresses in the range of (addresses) 10.10.20.0/24, translation port range 4000-5000, with a client connection limit of 100 connections per client. The translated address and port are persisted for 600 seconds. This object operates in NAT mode (Network Address and Port Translation mode), which is the default mode if not specified when type is dynamic-pat.

```
delete source-translation s1
```

Deletes the source-translation object named s1.

OPTIONS

app-service

Specifies the name of the application service to which this object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this object. Only the application service can modify or delete this object.

addresses

Specifies the set of translation IP addresses available in the pool. This is a collection of IP prefixes with their prefix lengths. All public-side addresses come from the subnets you enter in this property.

backup-addresses

Specifies translation IP addresses available in the backup pool which is used by DNAT translation mode if DNAT mode translation fails and falls back to NAT mode. This is a collection of IP prefixes with their prefix lengths.

client-connection-limit

The maximum number of simultaneous translated connections a client or subscriber is allowed to have. This attribute is applicable only if type is set to dynamic-pat.

description

User defined description.

egress-interfaces

The set of interfaces on which the source address translation is allowed or disallowed. If egress-interfaces-enabled is specified, the source address translation is allowed only on the specified set of interfaces. If egress-interfaces-disabled is specified, source address translation is disabled on specified interfaces.

egress-interfaces-disabled

Source address translation is not allowed on the interfaces specified in the egress-interfaces set.

egress-interfaces-enabled

Source address translation is allowed on the interfaces specified in the egress-interfaces set.

exclude-addresses

Specifies the set of addresses excluded from translation IP addresses available in the pool.

exclude-address-lists

Specifies the set of address lists (see security firewall address-list) excluded from translation IP addresses available in the pool.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

hairpin-mode

This attribute is applicable only if type is set to dynamic-pat.

Enable or disable hairpinning for incoming connections.

When a client sends a packet to another client in the same private network, hairpin mode sends the packet directly to the source client's private address; the BIG-IP system immediately translates the packet's public-side source address. Rather than going out to the public network and coming back later for translation, the packet takes a hairpin turn at the BIG-IP device.

icmp-echo

Enable or disable ICMP echo on translated addresses.

inbound-mode

This attribute is applicable only if type is set to dynamic-pat.

Modifies the inbound-connection mode for incoming connections to translation endpoints. A translation endpoint is the public-side address and port (X':x') for a private-side address (X:x). You can allow one of the following two algorithms for managing inbound connections:

endpoint-independent-filtering

creates inbound mappings automatically from outbound traffic and allows inbound connections.

Consider an outbound mapping from X:x to X':x'. If a connection comes from X:x through X':x', the BIG-IP system automatically creates a reverse mapping from X':x' back to X:x. A public-side station can respond through the X':x' address. This allows the BIG-IP system to provide Endpoint Independent Filtering (EIF) as defined in section 5 of RFC 4787
().

explicit

Allows inbound connections if and only if there exists an inbound mapping to translate public-side source address X':x' to client's private address X:x. Users can create Inbound mappings via iRules or PCP.

none disables inbound connections to translation end-points (X':x'). If there is a mapping of X (a private-side IP address) to X' (a public-side IP), connections can only go out from X through X'. If a public-side recipient tries to answer at the client's public-side X' address, the BIG-IP system does not map X' back to X. The inbound connection never happens.

Port Control Protocol (PCP) is not supported if you use this setting.

eif-timeout

Configurable range of eif-timeout is 3-300 seconds. Default value is 3 seconds. This attribute is only applicable if a) NAT method is Dynamic PAT (any pat-mode : NAT/PBA/DNAT) and b) inbound-mode is set to 'endpoint-independent-filtering'.

pat-mode

Specifies which kind of translation address mapping is performed when type is specified as dynamic-pat. Available options are NATP, Deterministic, and PBA.

NAPT (Network Address Port Translation) assigns translation addresses and ports in round-robin fashion. The algorithm first cycles through translation addresses and then through translation ports.

Deterministic

(DNAT) is a reversible translation method. A given client address and port always translates to a particular public address and port from the source-translation pool. This method has the following restrictions:

it is only available for NAT44 translations,
it does not support connections through DS-Lite tunnels,
subscriber connections must be received over a VLAN with the property, cmp-hash, set to "source ip,"
the egress to the Internet must be over a VLAN with the property, cmp-hash, set to "dest ip,"
any security NAT rule ("security nat policy") that uses this must have a source property set to an IP prefix containing fewer than 231 addresses. For example, the source cannot be 0.0.0.0/0.

PBA (Port Block Allocation) assigns 'blocks' of the translation addresses and ports to individual clients. All client connections are restricted to the allocated port blocks. Only block allocations and deallocations are logged in order to reduce the volume of logs.

subscriber connections must be received over a VLAN with the property, cmp-hash, set to "source ip,"
the egress to the Internet must be over a VLAN with the property, cmp-hash, set to "dest ip,"

You can access your VLAN configurations through the "net vlan" component. You can find the VLANs used by your virtual server by showing or listing the "itm virtual" component.

name Specifies a unique name for the source-translation component. This option is required for the commands create, delete, and modify.

mapping

These settings are applicable only if type is set to dynamic-pat.

Configure the mapping settings for translation entries. It is the preservation of a public-side IP address for a client from session to session.

mapping.mode

Configure the mapping mode for translation entries. You can enter address-pooling-paired, endpoint-independent-mapping, or none.

address-pooling-paired

causes the BIG IP software to attempt to keep the IP address persistent but not necessarily the port. If a client's private IP address:port combination is X:x, it's public-side address may be X':a in one session, X':b in the next session, X':c in a third session, and so on.

endpoint-independent-mapping

causes the BIG IP software to attempt to keep the IP address and port persistent. If a client's private IP address:port combination is X:x, and it's public-side address is X':x' in the first session, it remains X':x' in all future sessions.

This is called "Endpoint Independent Mapping" in RFC 4787 ().

This is the only supported setting for PCP, which you configure with the pcp property.

none prevents the BIG IP software from attempting any IP address or port mapping. An address:port combination of X:x is never guaranteed to have the same public-side address or port in two sessions.

mapping.timeout

After the most-recent session where address:port X:x translated to X':x' on the public side, a timer begins. If the timer expires before X:x has another session, X' or x' may be used as the public side of another address:port. Use this parameter to set the timeout (in seconds) for address and port mapping.

pcp A Port Control Protocol (PCP) client can set (or at least learn) its own translation (public-side) IP address and/or port. It can also set the address and/or port of a third-party client. PCP is defined in RFC 6887 (see).

pcp.profile

Specifies the PCP profile to use for this LSN pool. This PCP profile defines the settings to use for communication with PCP clients. Use the create itm profile pcp command to create a new PCP profile.

PCP requires a profile (defined with this property) and either a pcp.selfip or a pcp.dslite tunnel where clients can send their PCP requests.

If you remove this profile option, you must specifically remove any pcp.selfip or pcp.dslite tunnel, too.

pcp.selfip

Specifies the PCP Server self-IP address for this LSN pool. The virtual server's clients send their PCP packets to this address. Use the create net self command to create a self-IP address, then use that address for this parameter. Choose a self-IP address in a VLAN that is reachable by the virtual server's clients.

pcp.dslite

Specifies a DS-LITE tunnel for PCP packets. Whenever a client sends a PCP packet through this tunnel, the BIG-IP device uses the PCP profile you choose with the pcp.profile property.

A DS-LITE tunnel places each IPv4 packet into the payload of an IPv6 packet. The IPv6 packet carries the IPv4 packet between customer equipment and the BIG-IP system, which then removes the IPv4 packet, uses NAT to translate its IPv4 addresses, and sends it to its destination.

You cannot use this property if the `pat-mode` property is set to `Deterministic`.

`port-block-allocation`
Configures the port block settings for PBA mode.

`port-block-allocation.block-idle-timeout`
Configures the time after the last connection using the block is freed that the block assignment expires. The default value is 3600 seconds.

`port-block-allocation.block-lifetime`
Configures the timeout after which the block is no longer used for new port allocations. The block becomes a zombie block. The default is 0 which corresponds to an infinite timeout.

`port-block-allocation.block-size`
Configures the number of ports in a block. The default value is 64.

`port-block-allocation.client-block-limit`
Configures the number of blocks that can be assigned to a single subscriber IP address. The default value is 1.

`port-block-allocation.zombie-timeout`
Configures the timeout after which connections using the zombie block are killed. After connections are killed zombie block is freed after `port-block-allocation.block-idle-timeout`. This parameter is unused unless the `port-block-allocation.block-lifetime` is set. The default value is 0 which corresponds to infinite timeout.

`proxy-arp`
Enable or disable proxy arp for the translated source IP addresses.

`regex`
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

`ports`
Specifies the range of port numbers available for use with translation IP addresses.

`route-advertisement`
Enable or Disable Route Advertisement for the translated source IP addresses.

`type` Specifies which kind of translation is performed. Available options are: `static-nat`, `static-pat`, and `dynamic-pat`.

`dynamic-pat`
Using this type, BigIP translates a group of private (internal) IP Addresses to a pool of (one or more) public (external) IP Addresses and also translates ports to reuse the pool of public addresses. BigIP supports 3 different modes for `dynamic-pat` which can be specified using option `pat-mode` as described above.

`static-nat`
Using this translation type in the source-translation, only (source) IP Address translation is performed (and no port translation) for the incoming client traffic that matches the NAT Policy Rule using this translation object. There is a static (pre-defined) 1:1 mapping between the untranslated IP Address(es) and the translated IP Address(es) specified in this object (i.e same translation address (X') is used for all connections originating from the client with untranslated address (X)).

`static-pat`
Using this translation type in the source-translation, (source) port translation is performed for the incoming client traffic that matches the NAT Policy Rule using this translation object. There is a static (pre-defined) 1:1 mapping between the untranslated port(s) and the translated port(s) specified in this object. In addition, if translation address(es) are specified in the source-translation, it also performs IP Address translation (in the same fashion as done for `static-nat`).

SEE ALSO

`security nat policy`, `security nat destination-translation`, `ltm virtual`, `create`, `delete`, `edit`, `glob`, `list`, `ltm`, `modify`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2019-08-22 security nat source-translation(1)

security packet-filter default-rules

NAME

default-rules - Configures the default packet-filter rules. These rules are applied to packets if there's no more specific policy for these packets, for example, packets from Route Domain with an attached packet-filter policy will be processed with that Route Domain policy and not default packet-filter policy. Any particular packet is processed with at maximum one packet-filter policy.

MODULE

security packet-filter

SYNTAX

MODIFY

modify default-rules

options:

policy [[policy_name] | none]

edit default-rules

options:

all-properties

non-default-properties

DISPLAY

list default-rules

DESCRIPTION

You can use the default-rules component to configure packet-filter policy which is applied on all IPv6 traffic.

EXAMPLES

list default-rules

```
security packet-filter default-rules {  
  policy policy1  
}
```

Displays the current default policy.

OPTIONS

policy

Specifies an packet-filter policy. policy rules are applied by default.

SEE ALSO

edit, list, modify, security log profile, tmsh, security packet-filter policy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015-2018. All rights reserved.

BIG-IP 2018-05-10 security packet-filter default-rules(1)

security packet-filter policy

NAME

policy - Configures packet-filter policy.

MODULE

security packet-filter

SYNTAX

Modify the policy component within the security packet-filter module using the syntax shown in the following sections.

CREATE/MODIFY

create policy [name]

modify policy [name]

options:

description [string]

rules [add | delete | modify | replace-all-with] {

[name] {

options:

action [accept | drop]

description [string]

ipv6-extension-headers [add | delete | replace-all-with] {

[ah | esp | hopopt | ipv6-frag | ipv6-nonxt | ipv6-opts | ipv6-route | mh] {

```
values [add | delete | replace-all-with] {
  [ [value] | [value1-value2] ]
}
}
log [no | yes]
status [disabled | enabled]
}
}
rules none
```

```
edit policy
options:
  all-properties
  non-default-properties
```

```
DISPLAY
list policy
```

DESCRIPTION

You can use the policy component to configure a shareable and reusable set of security packet-filter rules which can be associated with a number of configuration objects of the following types: security packet-filter default-rules, net route-domain.

EXAMPLES

```
modify policy policy1 rules add {
drop-frags {
  action drop
  ipv6-extension-headers replace-all-with { ipv6-frag }
} }
```

Creates a rule entry that drops all IPv6 packets specifying Fragment Header.

```
modify policy policy1 rules delete drop-frags
```

Removes the rule drop-frags from the list of rules.

```
create policy xyz rules add { r1 { action drop ipv6-extension-headers replace-all-with { hopopt { values
replace-all-with { 0 } } } } }
```

Creates a policy with a single rule that drops all packets specifying Hop-by-Hop Header with option 0. Packets specifying Hop-by-Hop Header, but without option 0, will not be dropped.

```
create policy xyz rules add { r1 { action drop ipv6-extension-headers replace-all-with { hopopt } } }
```

Creates a policy with a single rule that drops all packets specifying Hop-by-Hop Header with or without any options.

```
list policy
```

Displays the current list of policy rules.

OPTIONS

```
description
User defined description.
```

```
rules
Adds, deletes, or replaces a packet-filter rule.
```

```
action
Specifies the action that the system takes when a rule is matched.
```

```
accept
Specifies that the current packet should be accepted.
```

```
drop Specifies that the current packet should be silently dropped. Nothing is sent back to the
packet source. The packet is still compared to any other rules, so other rule counters may be
incremented.
```

```
description
User defined description.
```

```
ipv6-extension-headers
Specifies a list of IPv6 Extension Header types (only one item per list is currently supported),
against which the packet will be compared
```

```
values
Specifies a list of IPv6 Extension Header options or option ranges against which the packet
will be compared. Specifying values is supported only for some of the IPv6 Extension Header
types: hopopt - values match Hop-by-Hop Options; ipv6-opts - values match Destination Options;
ipv6-route - values match Routing type. If values are omitted - no values are required in
packet to match the rule, so it's enough for a packet to include the IPv6 Extension Header of
the type specified in a rule for rule to match.
```

```
log Specifies whether the packet will be logged if it matches the rule. Logging must also be enabled in
security log profile global-network. Note that the statistics counter is always incremented when a
packet matches a rule.
```

status
Specifies whether the rule is enabled or disabled. If a rule is disabled it has no effect on data packets.

SEE ALSO
create, edit, list, modify, security log profile, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2018. All rights reserved.

BIG-IP 2018-05-10 security packet-filter policy(1)

security packet-filter rule-stat

NAME
rule-stat - Displays statistics of packet-filter rules on the BIG-IP(r) system. You can only use the show command with this component.

MODULE
security packet-filter

SYNTAX
show rule-stat
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
field-fmt

DESCRIPTION
You can use the rule-stat component to display statistics of packet-filter rules.

EXAMPLES
show rule-stat

Displays packet-filter rule's statistics in the system default units.

show rule-stat raw

Displays raw packet-filter rule's statistics.

SEE ALSO
show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2018. All rights reserved.

BIG-IP 2018-05-10 security packet-filter rule-stat(1)

security presentation tmui netflow-details

NAME
netflow-details - List Specified netflow protected server.

MODULE
security presentation tmui

SYNTAX
Configure the netflow-details component within the security presentation tmui module using the syntax shown in the following sections.

DISPLAY
list netflow-details [name]

DESCRIPTION

You can use the netflow-details component to display the specified netflow protected server configuration and stats info.

EXAMPLES

```
list netflow-details my_netflow_protected_server
```

Displays the properties of the named my_netflow_protected_server configuration and stats info.

OPTIONS

name Specifies the netflow protected server name.

SEE ALSO

list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-31 security presentation tmui netflow-details(1)

security presentation tmui netflow-list

NAME

netflow-list - List ALL Netflow Protected Server.

MODULE

security presentation tmui

SYNTAX

Configure the netflow-list component within the security presentation tmui module using the syntax shown in the following sections.

DISPLAY

```
list netflow-list
```

DESCRIPTION

You can use the netflow-list component to display all the current netflow protected server configuration and stats info.

EXAMPLES

```
list netflow-list
```

Displays the properties of all the current netflow protected server configuration and stats info.

SEE ALSO

list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-31 security presentation tmui netflow-list(1)

security presentation tmui signature-details

NAME

signature-details - List Specified DoS Dynamic or Persistent Signature.

MODULE

security presentation tmui

SYNTAX

Configure the signature-details component within the security presentation tmui module using the syntax shown in the following sections.

DISPLAY

list signature-details [name]

DESCRIPTION

You can use the signature-details component to display the specified dos dynamic or persistent signature configuration and stats info.

EXAMPLES

```
list signature-details my_signature
```

Displays the properties of the named my_signature configuration and stats info.

OPTIONS

name Specifies the signature name.

SEE ALSO

list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-31 security presentation tmui signature-details(1)

security presentation tmui signature-list

NAME

signature-list - List ALL DoS Dynamic and Persistent Signatures.

MODULE

security presentation tmui

SYNTAX

Configure the signature-list component within the security presentation tmui module using the syntax shown in the following sections.

DISPLAY

```
list signature-list
```

DESCRIPTION

You can use the signature-list component to display all the current dos dynamic and persistent signatures configuration and stats info.

EXAMPLES

```
list signature-list
```

Displays the properties of all the current dos dynamic and persistent signatures configuration and stats info.

SEE ALSO

list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-31 security presentation tmui signature-list(1)

security protected-servers netflow-tmc-stat

NAME

netflow-tmc-stat - Show netflow protected server's related stats.

MODULE

security protected-servers

SYNTAX

```
show netflow-tmc-stat
```

DESCRIPTION

You can use the netflow-tmc-stat component to display Traffic-Matching-Criteria, megabits-per-second, kilo-packets-per-second, and "connections per second.

EXAMPLES

```
show security protected-servers netflow-tmc-stat
```

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 security protected-servers netflow-tmc-stat(1)

security protocol-inspection auto-update settings

NAME

auto-update settings - Display the protocol inspection auto update settings.

MODULE

security protocol-inspection auto-update settings

SYNTAX

LIST

Options:

|

Properties:

auto-update-interval partition

enabled {

DISPLAY

```
list security protocol-inspection auto-update settings
```

DESCRIPTION

Use this command to list the settings of protocol inspection auto-update for automatic hitless upgrade of im package.

OPTIONS

For information about the options that you can use with the command list, see help list.

EXAMPLES

```
list security protocol-inspection auto-update settings
```

Display auto update property.

SEE ALSO

B, B, B

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection auto-update settings(1)

security protocol-inspection auto-update status

NAME

auto-update status - Display the protocol inspection auto update status.

MODULE

security protocol-inspection auto-update status

SYNTAX

LIST

Options:

|

Properties:

last-updated-time partition {
message progress-status

DISPLAY

list security protocol-inspection auto-update status

DESCRIPTION

Use this command to list the status of protocol inspection auto-update for automatic hitless upgrade of im package.

OPTIONS

For information about the options that you can use with the command list, see help list.

EXAMPLES

list security protocol-inspection auto-update status

Display auto update property.

SEE ALSO

B, B, B

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection auto-update status(1)

security protocol-inspection common-config

NAME

common-config - Configures the protocol inspection common-configs.

MODULE

security protocol-inspection common-config

SYNTAX

CREATE/MODIFY

modify security protocol-inspection common-config

create security protocol-inspection common-config

properties:

app-service [string]

compliance { ... }

description [string]

service { ... }

DISPLAY

list security protocol-inspection common-config

DESCRIPTION

Use this command to create/modify protocol inspection common-config.

EXAMPLES

create security protocol-inspection common-config new_common_config

Create common-config "new_common_config".

list security protocol-inspection common-config

Displays all protocol inspection common-configs.

list security protocol-inspection common-config new_common_config

Displays protocol inspection common-config with name new_common_config.

PROPERTIES

app-service [string]

Specifies app service.

compliance

Specifies common-config compliances for this common-config. ...

description [string]

Specifies the description.

service

Specifies common-config services for this common-config. ...

SEE ALSO

create, modify, list, security, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2019. All rights reserved.

BIG-IP 2019-07-25 security protocol-inspection common-config(1)

security protocol-inspection compliance-enums

NAME

compliance enums - Show available compliance enums.

MODULE

security protocol-inspection compliance-enums

SYNTAX

DISPLAY

list security protocol-inspection compliance-enums

properties:

app-service [string]

description [string]

insp-id [integer]

value [string]

DESCRIPTION

Use this command to get information compliance check enums. This command does not allow you to create/modify compliance check enums.

EXAMPLES

list security protocol-inspection compliance-enums

Displays all available compliance enums.

PROPERTIES

app-service

Specifies app service.

description

Specifies description.

insp-id

Specifies compliance check identifier.

value

Specifies compliance check value.

SEE ALSO

list, security, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017. All rights reserved.

BIG-IP 2017-11-21 security protocol-inspection compliance-enums(1)

security protocol-inspection compliance

NAME

compliance - Show compliance inspections (this section is not modifiable)

MODULE

security protocol-inspection compliance

SYNTAX

DISPLAY

list security protocol-inspection compliance

options:

all

all-properties

non-default-properties

one-line

properties:

accuracy [high | low | medium]

description [string]

service [string]

action [accept | drop | reject]

direction [any | to-client | to-server]

log [yes | no]

app-service [string]

documentation [string]

performance-impact [high | low | medium]

systems [string]

attack-type [string]

id [integer]

protocol [any | tcp | udp]

risk [critical | high | low | medium]

value [string]

value-type [int | string | bool | enum | vector-string | vector-int | vector-enum]

DISPLAY

list security protocol-inspection compliance

DESCRIPTION

Use this command to get default information about compliance checks. This command does not allow you to create/modify compliance checks.

EXAMPLES

list security protocol-inspection compliance

Displays all compliance inspections.

list security protocol-inspection compliance ftp_malformed_param

Displays compliance inspection "ftp_malformed_param".

PROPERTIES

accuracy

Specifies the accuracy of the compliance inspection.

description

Specifies the description of the compliance inspection. Also this parameter is used in logging when compliance inspection is matched.

service

Specifies target-based service.

action

Specifies enforcement action for matched compliance inspection.

direction

Specifies the flow direction for which this compliance inspection will be applied.

log Specifies whether the inspection will be logged if it matches the compliance inspection.

app-service

Specifies app service.

documentation

Specifies compliance inspection documentation.

performance-impact

Specifies performance impact of this compliance inspection.

systems

Specifies systems where this compliance inspection can be matched.

attack-type

Specifies compliance inspection attack type.

id Specifies compliance inspection identifier.

protocol

Specifies transport protocol where this compliance inspection can be matched (udp, tcp, any).

risk Specifies compliance inspection risk.

value

Specifies value for specific compliance check depending on field (see below).

value-type
Specifies value type for specific compliance check.

SEE ALSO

list, security, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017. All rights reserved.

BIG-IP 2018-01-10 security protocol-inspection compliance(1)

security protocol-inspection learning-stats

NAME

learning-stats - Delete the learning stats of protocol inspection.

MODULE

security protocol-inspection learning-stats

SYNTAX

DELETE

delete security protocol-inspection learning-stats

Options:

all

DISPLAY

delete security protocol-inspection learning-stats

DESCRIPTION

Use this command to delete the learning stats of protocol inspection.

OPTIONS

For information about the options that you can use with the command delete, see help delete.

EXAMPLES

delete security protocol-inspection learning-stats

Delete the learning stats.

SEE ALSO

B, B, B

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection learning-stats(1)

security protocol-inspection learning-suggestions

NAME

learning-suggestions - Display the suggestions learned from the protocol inspection.

MODULE

security protocol-inspection learning-suggestions

SYNTAX

SHOW

show security protocol-inspection learning-suggestions

Options:

field-fmt include-reason

include-published |

DISPLAY

show security protocol-inspection learning-suggestions

DESCRIPTION

Use this command to display the learned suggestions of the protocol inspection (suggested actions of inspections learned from the past traffic pattern that hit the inspections).

OPTIONS

For information about the options that you can use with the command show, see help show.

EXAMPLES

```
show security protocol-inspection learning-suggestions
```

Display the suggested learnings.

SEE ALSO

B, B, B

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection learning-suggestions(1)

security protocol-inspection profile-status

NAME

learning-suggestions - Display the status of protocol inspection profiles.

MODULE

security protocol-inspection profile-status

SYNTAX

SHOW

```
show security protocol-inspection profile-status
```

Options:

|

DISPLAY

```
show security protocol-inspection profile-status
```

DESCRIPTION

Use this command to display the status (ready or not due to signature blob compilation) of the protocol inspection profiles.

OPTIONS

For information about the options that you can use with the command show, see help show.

EXAMPLES

```
show security protocol-inspection profile-status
```

Display the profile status if its ready or not.

SEE ALSO

B, B, B

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection profile-status(1)

security protocol-inspection profile

NAME

profile - Configures the protocol inspection profiles.

MODULE

security protocol-inspection profile

SYNTAX

CREATE/MODIFY

modify security protocol-inspection profile

create security protocol-inspection profile

properties:

app-service [string]
auto-add-new-inspections [bool]
auto-publish-suggestion [bool]
avr-stat-collect [bool]
common-config [string]
common-config-merge-type [string]
compliance-enable [bool]
defaults-from [string]
description [string]
signature-enable [bool]
services { ... }
staging-period [integer]

DISPLAY

list security protocol-inspection profile

DESCRIPTION

Use this command to create/modify protocol inspection profile.

EXAMPLES

```
create security protocol-inspection profile new_http_profile { default-from protocol_inspection_http }
```

Create profile "new_http_profile" and clone all configuration from predefined profile protocol_inspection_http.

```
list security protocol-inspection profile
```

Displays all protocol inspection profiles.

```
list security protocol-inspection profile new_http_profile
```

Displays protocol inspection profile with name new_http_profile.

```
modify security protocol-inspection profile new_http_profile common-config common_config_name
```

Attach common-config to the profile.

PROPERTIES

app-service [string]

Specifies app service.

auto-add-new-inspections [bool]

If set, new inspections arrived via IPS IM pkg will be automatically added to this profile for configured service.

auto-publish-suggestion [bool]

If set, after the learning period(staging-period) the action for the inspections will automatically be updated to the suggested action, if exists.

avr-stat-collect [bool]

Specifies if AVR collects data from IPS.

common-config [string]

Specifies common-config to be attached to the profile.

common-config-merge-type [string]

Specifies common-config value merge type.

compliance-enable [bool]

Specifies whether the compliance checks will be enabled for this profile.

signature-enable [bool]

Specifies whether the signature checks will be enabled for this profile.

defaults-from [string]

Specifies parent profile (in time of creating). If this parameter is assigned then new profile will be cloned from parent profile.

description [string]

Specifies profile description.

services

Specifies services for this profile. ...

staging-period

If auto-publish-suggestion is set, this value defines the time period after which inspection action will be automatically updated to suggested action.

SEE ALSO

create, modify, list, security, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017. All rights reserved.

BIG-IP 2019-07-25 security protocol-inspection profile(1)

security protocol-inspection service

NAME

service - Show services (this section is not modifiable)

MODULE

security protocol-inspection service

SYNTAX

CREATE/MODIFY Creating/modifying services are disallowed.

DISPLAY

list security protocol-inspection service

options:

all

all-properties

non-default-properties

one-line

properties:

app-service [string]

description [string]

id [integer]

partition [string]

ports [string]

DISPLAY

list security protocol-inspection service

DESCRIPTION

Use this command to get information about services. This command does not allow you to create/modify services.

EXAMPLES

list security protocol-inspection service

Displays all services.

list security protocol-inspection signature dns

Displays service "dns".

PROPERTIES

app-service

Specifies app service.

description

Specifies service description.

id Specifies service identifier.

partition

Specifies service partition.

ports

Specifies default service ports in format "80, 81".

SEE ALSO

list, security, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017. All rights reserved.

BIG-IP 2017-11-20 security protocol-inspection service(1)

security protocol-inspection signature

NAME

signature - Configures the signature inspections.

MODULE

security protocol-inspection signature

SYNTAX

```
CREATE/MODIFY
modify security protocol-inspection signature
create security protocol-inspection signature
properties:
accuracy [high | low | medium]
description [string]
last-updated [date in format %y-%m-%d:%H:%M:%S]
reference-links [string]
service [string]
action [accept | drop | reject]
direction [any | to-client | to-server]
log [yes | no]
references [string]
sig [string - signature in snort format]
app-service [string]
documentation [string]
performance-impact [high | low | medium]
revision [integer]
systems [string]
attack-type [string]
id [integer]
protocol [any | tcp | udp]
risk [critical | high | low | medium]
user-defined [yes | no]
```

DISPLAY

```
list security protocol-inspection signature
```

DESCRIPTION

Use this command to create/modify custom signatures in snort format.

EXAMPLES

```
create security protocol-inspection signature new_sig { log yes action drop sig "content:\\"GET\\";
content:\\"HTTP\\";" description "Signature match" }
```

Create signature "new_sig" which find "GET" and "HTTP" in payload (see details about snort signatures in related documentation). Following actions are applied if signature is matched: drop flow and write message "Signature match".

```
modify security protocol-inspection signature new_sig { log no action accept sig }
```

Modify action and logging of previous signature "new_sig". Following actions are applied if signature is matched: accept flow.

```
list security protocol-inspection signature new_sig
```

Displays signature new_sig.

```
list security protocol-inspection signature
```

Displays all signatures.

PROPERTIES

accuracy
Specifies the accuracy of the signature.

description
Specifies the description of the signature. Also this parameter is used in logging when signature is matched.

last-updated
Specifies date/time when signature has been updated last time.

reference-links
Specifies external references (url) to signature.

references
Specifies external industrial references (cve and bugtraq) to signature.

service
Specifies target-based service.

action

Specifies enforcement action for matched signature.

direction

Specifies flow direction for signature. Signature search will apply only for payload in this direction.

log Specifies whether the inspection will be logged if it matches the signature.

app-service

Specifies app service.

documentation

Specifies signature documentation.

performance-impact

Specifies performance impact of this signature.

revision

Specifies signature revision. For custom signatures, this parameter will be incremented each time you modify this signature.

systems

Specifies systems where this signature can be matched.

attack-type

Specifies signature attack type.

id Specifies signature identifiers.

protocol

Specifies transport protocol where this signature can be matched (udp, tcp, any).

risk Specifies signature risk.

sig Specifies snort signature.

user-defined

Specifies if signature is created by user.

deprecated

Specifies if inspection is now deprecated and will not be matched anymore.

SEE ALSO

list, modify, security, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017. All rights reserved.

BIG-IP 2018-01-11 security protocol-inspection signature(1)

security protocol-inspection staging

NAME

staging - Display the staged inspections of protocol inspection.

MODULE

security protocol-inspection staging

SYNTAX

SHOW

show security protocol-inspection staging

Options:

field-fmt |

DISPLAY

show security protocol-inspection staging

DESCRIPTION

Use this command to display the staging properties of inspections.

OPTIONS

For information about the options that you can use with the command show, see help show.

EXAMPLES

show security protocol-inspection staging

Display the staged inspections.

SEE ALSO

B, B, B

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection staging(1)

security protocol-inspection system

NAME

system - Display the system compliances from the protocol inspection.

MODULE

security protocol-inspection system

SYNTAX

LIST

list security protocol-inspection system

Options:

all
all-properties
non-default-properties
one-line
|

Properties:

accuracy
action
app-service
attack-type
deprecated
description
documentation
id
log
partition
performance-impact
risk
value
value-type
{

DISPLAY

list security protocol-inspection system

DESCRIPTION

Use this command to display the system compliances of protocol inspection.

OPTIONS

For information about the options that you can use with the command list, see help list.

EXAMPLES

list security protocol-inspection system

Display the system compliances.

SEE ALSO

B, B, B

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection system(1)

security protocol-inspection updates

NAME
updates - Configures updates.

MODULE
security protocol-inspection updates

SYNTAX
INSTALL
install security protocol-inspection updates
options:
file [string]

DISPLAY
show security protocol-inspection updates

DESCRIPTION
Use this command to install and see all available IM packages.

EXAMPLES
install security protocol-inspection updates file file_name

Install new update from file "file_name".

show security protocol-inspection updates

Show all available updates.

OPTIONS
file Specifies IM package file.

SEE ALSO
show, install, security, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017. All rights reserved.

BIG-IP 2017-11-21 security protocol-inspection updates(1)

security protocol-inspection virtual-servers

NAME
virtual-servers - Display the virtual servers attached to the protocol inspection profile.

MODULE
security protocol-inspection virtual-servers

SYNTAX
SHOW
show security protocol-inspection virtual-servers
Options:
field-fmt |
Profiles:

DISPLAY
show security protocol-inspection virtual-servers

DESCRIPTION
Use this command to display the virtual servers attached to protocol inspection profile.

OPTIONS
For information about the options that you can use with the command show, see help show.

EXAMPLES
show security protocol-inspection virtual-servers protocol-inspection

Display the virtual servers attached to the profile.

SEE ALSO
B, B, B

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2017, 2019. All rights reserved.

BIG-IP 2019-04-25 security protocol-inspection virtual-servers(1)

security scrubber dwbl-scrubber-category-stats

NAME

dwbl-scrubber-category-stats - Show the list of IPs that are associated with a scrubbed category.

MODULE

security scrubber

SYNTAX

show dwbl-scrubber-category-stats

DESCRIPTION

You can use the dwbl-scrubber-category-stats component to display the list of IPs that are associated with a scrubbed category.

EXAMPLES

show security scrubber dwbl-scrubber-category-stats category-name [name] profile-name scrubber-profile-default

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2016-07-07 security scrubber dwbl-scrubber-category-stats(1)

security scrubber dwbl-scrubber-stat

NAME

dwbl-scrubber-stat - Show scrubber's Dwbl related stats.

MODULE

security scrubber

SYNTAX

Run the show dwbl-scrubber-stat within the security scrubber module using the syntax shown in the following sections.

SHOW

show dwbl-scrubber-stat

options:

destination
expire-time
monitors
observed-rate-bps
observed-rate-cps
observed-rate-pps
route-domain
scrubber-status
scrubber-type
silverline-status
silverline-url
threshold-bps
threshold-cps
threshold-pps
ttl
under-attack

DESCRIPTION

You can use the dwbl-scrubber-stat component to display the threshold, observed rate, under attack, time to live, and expire time for each monitor. Each row displays stats related per a monitor. Scrubber module monitors selected virtual servers, categories, route domains, and netflow protected servers.

EXAMPLES

show security scrubber dwbl-scrubber-stat show security scrubber dwbl-scrubber-stat monitors show security scrubber dwbl-scrubber-stat scrubber-status show security scrubber dwbl-scrubber-stat silverline-url silverline-status expire-time under-attack

OPTIONS

destination

Show the column for the destination configured on the scrubber.

expire-time

Show the column for the current expiration time on the scrubber.

monitors

Show the columns related to the monitor configuration: threshold-bps, observed-rate-bps, threshold-cps, observed-rate-cps, threshold-pps, observed-rate-pps.

observed-rate-bps

Show the column for the current observed BPS on the monitor.

observed-rate-cps

Show the column for the current observed CPS on the monitor.

observed-rate-pps

Show the column for the current observed PPS on the monitor.

route-domain

Show the column for the Route Domain configured on the monitor.

scrubber-status

Show the columns related to the current status of the scrubber: under-attack, expire-time, silverline-status.

scrubber-type

Show the column for the type of monitor configured for the scrubber.

silverline-status

Show the column for the current status of Silverline.

silverline-url

Show the column for the URL configured for Silverline.

threshold-bps

Show the column for the BPS threshold configured for the scrubber.

threshold-cps

Show the column for the CPS threshold configured for the scrubber.

threshold-pps

Show the column for the PPS threshold configured for the scrubber.

ttd Show the column for the TTL configured for the scrubber.

under-attack

Show the column for whether an attack has been detected.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2018-01-10 security scrubber dwbl-scrubber-stat(1)

security scrubber profile

NAME

profile - Configures a scrubber profile for use by firewall. A scrubber-profile-default specifies monitors and method (how and where) to be monitored and scrubbed.

MODULE

security scrubber

SYNTAX

Configure the scrubber-profile-default component within the security scrubber profile module using the syntax in the following sections.

MODIFY

modify profile [name]
options:

```

advertisement-ttl [integer]
scrubber-categories action [add | delete | modify | none | replace-all-with] {
  [name] {
    options:
advertisement-method [bgp-flowspec-method | bgp-method | none-method | silverline-method]
app-service [[string] | none]
bgp-flowspec-advertisement-action [drop | redirect | rate-limit | qos]
bgp-flowspec-dscp-value [integer]
bgp-flowspec-rate-limit [integer]
bgp-flowspec-redirect-asn-community [string]
blacklist-category [string]
next-hop [IPv4 address]
next-hop-v6 [IPv6 address]
route-domain-name [string]
  }
}
scrubber-netflow-protected-server [add | delete | modify | none | replace-all-with] {
  [name] {
    options:
advertisement-method [bgp-flowspec-method | bgp-method | none-method | silverline-method]
app-service [[string] | none]
bgp-flowspec-advertisement-action [drop | redirect | rate-limit | qos]
bgp-flowspec-dscp-value [integer]
bgp-flowspec-rate-limit [integer]
bgp-flowspec-redirect-asn-community [string]
blacklist-category [string]
next-hop [IPv4 address]
next-hop-v6 [IPv6 address]
route-domain-name [string]
  }
}
scrubber-rt-domain action [add | delete | modify | none | replace-all-with] {
  [name] {
    options:
absolute-threshold [integer]
advertisement-method [bgp-flowspec-method | bgp-method | none-method | silverline-method]
bgp-flowspec-advertisement-action [drop | redirect | rate-limit | qos]
bgp-flowspec-dscp-value [integer]
bgp-flowspec-rate-limit [integer]
bgp-flowspec-redirect-asn-community [string]
next-hop [IPv4 address]
next-hop-v6 [IPv6 address]
percentage-threshold [integer]
route-domain [string]
scrubber-rd-network-prefix action [add | delete | modify | none | replace-all-with] {
  [name] {
    options:
app-service [[string] | none]
bgp-flowspec-advertisement-action [drop | redirect | rate-limit | qos]
bgp-flowspec-dscp-value [integer]
bgp-flowspec-rate-limit [integer]
bgp-flowspec-redirect-asn-community [string]
dst-ip [IP address]
mask [integer]
next-hop [IP address]
  }
}
excluded-vlans action [add | delete | none | replace-all-with] {
  [name] {}
}
}
scrubber-virtual-server action [add | delete | modify | none | replace-all-with] {
  [name] {
    options:
absolute-threshold [integer]
advertisement-method [bgp-flowspec-method | bgp-method | none-method | silverline-method]
app-service [[string] | none]
bgp-flowspec-advertisement-action [drop | redirect | rate-limit | qos]
bgp-flowspec-dscp-value [integer]
bgp-flowspec-rate-limit [integer]
bgp-flowspec-redirect-asn-community [string]
next-hop [IPv4 address]
next-hop-v6 [IPv6 address]
percentage-threshold [integer]
vs-name [string]
  }
}
silverline { url [string] user-id [string] user-passwd [string] }
app-service [[string] | none]

list profile [[name] | all | [property]]
show running-config profile [[name] | all | [property]]
options:
all-properties
non-default-properties
one-line

```

recursive

OPTIONS

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

description

User defined description.

advertisement-ttl

Defines the scrubbing duration for all monitored entities in seconds.

scrubber-categories

Defines how a blacklist-category to be scrubbed.

OPTIONS

advertisement-method

Defines a method to use to scrub a blacklist-category.

bgp-flowspec-advertisement-action

Specifies the BGP FlowSpec Advertisement Action to be used for scrubbing Blacklist category. The default is redirect

bgp-flowspec-dscp-value

Specifies the BGP FlowSpec DSCP value for advertisement qos action.

bgp-flowspec-rate-limit

Specifies the BGP FlowSpec rate limit (bytes/sec) for advertisement rate limiting action.

bgp-flowspec-redirect-asn-community

Specifies the BGP Extended Community value (in the format - AA:NNN, where AA is 16-bit number and NNN is 32-bit number) for redirect-to-VRF support when BGP Flowspec advertisement action is redirect.

blacklist-category

Identifies a blacklist-category to be scrubbed.

next-hop

Defines the nexthop to be used for scrubbing/redirecting traffic for IPv4 shuns.

next-hop-v6

Defines the nexthop to be used for scrubbing/redirecting traffic for IPv6 shuns.

route-domain-name

Identifies a route-domain to be used for route advertisement.

OPTIONS

absolute-threshold

Specifies aggregate maximum bandwidth threshold in Mbps.

advertisement-method

Defines a method to use to scrub a NetFlow protected server object.

app-service

The application service that the object belongs to.

bgp-flowspec-advertisement-action

Specifies the BGP FlowSpec Advertisement Action to be used for scrubbing NetFlow protected server. The default is redirect.

bgp-flowspec-dscp-value

Specifies the BGP FlowSpec DSCP value for advertisement qos action.

bgp-flowspec-rate-limit

Specifies the BGP FlowSpec rate limit (bytes/sec) for advertisement rate limiting action.

bgp-flowspec-redirect-asn-community

Specifies the BGP Extended Community value (in the format - AA:NNN, where AA is 16-bit number and NNN is 32-bit number) for redirect-to-VRF support when BGP Flowspec advertisement action is redirect.

cps-absolute-threshold

Specifies aggregate maximum connection threshold in CPS (Connection Per Second).

cps-percentage-threshold

Specifies aggregate maximum connection rate (CPS) threshold as a percentage of NetFlow capacity.

next-hop

Specifies BGP redirection next hop property.

nps-name

Specifies the name of the specified NetFlow protected server.

percentage-threshold

Specifies aggregate maximum bandwidth (BPS) threshold as a percentage of NetFlow capacity.

pps-absolute-threshold

Specifies aggregate maximum packet threshold in PPS (Packet Per Second).

pps-percentage-threshold

Specifies aggregate maximum packet rate (PPS) threshold as a percentage of NetFlow capacity.

OPTIONS

absolute-threshold

Defines bandwidth threshold which triggers scrubbing for selected route domain.

advertisement-method

Defines a method to use to scrub a route domain.

bgp-flowspec-advertisement-action

Specifies the BGP FlowSpec Advertisement Action to be used for scrubbing a route domain. The default is redirect.

bgp-flowspec-dscp-value

Specifies the BGP FlowSpec DSCP value for advertisement qos action.

bgp-flowspec-rate-limit

Specifies the BGP FlowSpec rate limit (bytes/sec) for advertisement rate limiting action.

bgp-flowspec-redirect-asn-community

Specifies the BGP Extended Community value (in the format - AA:NNN, where AA is 16-bit number and NNN is 32-bit number) for redirect-to-VRF support when BGP Flowspec advertisement action is redirect.

percentage-threshold

Defines bandwidth threshold which triggers scrubbing for selected route domain. The percentage is calculate based on route-domain bandwidth value.

next-hop

Defines the nexthop to be used for scrubbing/redirecting IPv4 traffic.

next-hop-v6

Defines the nexthop to be used for scrubbing/redirecting IPv6 traffic.

route-domain-name

Identifies a route-domain to be used for route advertisement.

excluded-vlans

Identifies VLANs to be excluded from traffic monitoring.

scrubber-rd-network-prefix

Defines subnets which to be used for scrubbing/redirecting traffic. If is defined than the scrubbing for parent route-domain would be ignored.

OPTIONS

bgp-flowspec-advertisement-action

Specifies the BGP FlowSpec Advertisement Action to be used for scrubbing route domain subnets. The default is redirect.

bgp-flowspec-dscp-value

Specifies the BGP FlowSpec DSCP value for advertisement qos action.

bgp-flowspec-rate-limit

Specifies the BGP FlowSpec rate limit (bytes/sec) for advertisement rate limiting action.

bgp-flowspec-redirect-asn-community

Specifies the BGP Extended Community value (in the format - AA:NNN, where AA is 16-bit number and NNN is 32-bit number) for redirect-to-VRF support when BGP Flowspec advertisement action is redirect.

dst-ip

Defines subnet to be used for redirection.

mask

Defines subnet mask to be used for redirection.

next-hop

Defines the nexthop to be used for scrubbing/redirecting traffic.

app-service

Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

scrubber-virtual-server

Defines how and when a virtual server to be scrubbed.

OPTIONS

absolute-threshold

Defines a bandwidth threshold which triggers scrubbing for a selected virtual server.

advertisement-method

Defines a method to use to scrub a virtual server.

bgp-flowspec-advertisement-action

Specifies the BGP FlowSpec Advertisement Action to be used for scrubbing a virtual server. The default is redirect.

`bgp-flowspec-dscp-value`
Specifies the BGP FlowSpec DSCP value for advertisement qos action.

`bgp-flowspec-rate-limit`
Specifies the BGP FlowSpec rate limit (bytes/sec) for advertisement rate limiting action.

`bgp-flowspec-redirect-asn-community`
Specifies the BGP Extended Community value (in the format - AA:NNN, where AA is 16-bit number and NNN is 32-bit number) for redirect-to-VRF support when BGP Flowspec advertisement action is redirect.

`percentage-threshold`
Defines bandwidth threshold which triggers scrubbing for selected route domain. The percentage is calculate based on defined virtual server bandwidth value.

`next-hop`
Defines the nexthop to be used for scrubbing/redirection traffic for IPv4 VS destination addresses.

`next-hop-v6`
Defines the nexthop to be used for scrubbing/redirection traffic for IPv6 VS destination addresses.

`vs-name`
Identifies a virtual server to be used for route advertisement.

`app-service`
Specifies the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the Application Service that owns the object, you cannot modify or delete the object. Only the Application Service can modify or delete the object.

OPTIONS

`url`
Used to communicate with Silverline system.

`user-id`
Defines silverline user's user identification.

`user-passwd`
Defines silverline user's password.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2019-12-09 security scrubber profile(1)

security scrubber unredirect

NAME

`unredirect` - Stop scrubbing for a monitored entity. A `unredirect` command enables a user to stop redirection(scrubbing) for a selected monitored entity such as category, virtual server, or route domain.

MODULE

security scrubber

SYNTAX

Configure the redirect component within the security scrubber module using the syntax shown in the following sections.

RUN

```
run security scrubber unredirect
options:
  profile [scrubber-profile-default]
  unredirect-category [string]
  unredirect-netflow-protected-server [string]
  unredirect-route-domain [string]
  unredirect-virtual-server [string]
```

DESCRIPTION

You can use the `unredirect` component to stop traffic redirection (scrubbing) for a selected monitored entity.

EXAMPLES

```
run security scrubber unredirect scrubber-profile scrubber-profile-default unredirect-category
run security scrubber unredirect scrubber-profile scrubber-profile-default unredirect-virtual-server
run security scrubber unredirect scrubber-profile scrubber-profile-default unredirect-route-domain
```

run security scrubber unredirect scrubber-profile scrubber-profile-default unredirect-netflow-protected-server

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2017-06-20 security scrubber unredirect(1)

security ssh profile

NAME

profile - Configures ssh profile.

MODULE

security ssh

SYNTAX

Modify the profile component within the security ssh module using the syntax shown in the following sections.

CREATE/MODIFY

```
create profile [name]
modify profile [name]
options:
  description [string]
  rules [add | delete | modify | replace-all-with] {
    [ [name] ] {
      options:
    }
  }
actions [add | delete | modify] {
  [ [name] ] {
    shell-action { control [allow | disallow | terminate] log [no | yes] }
    sub-system-action { control [allow | disallow | terminate] log [no | yes] }
    sftp-up-action { control [allow | disallow | terminate] log [no | yes] }
    sftp-down-action { control [allow | disallow | terminate] log [no | yes] }
    scp-up-action { control [allow | disallow | terminate] log [no | yes] }
    scp-down-action { control [allow | disallow | terminate] log [no | yes] }
    rexec-action { control [allow | disallow | terminate] log [no | yes] }
    local-forward-action { control [allow | disallow | terminate] log [no | yes] }
    remote-forward-action { control [allow | disallow | terminate] log [no | yes] }
    x11-forward-action { control [allow | disallow | terminate] log [no | yes] }
    agent-action { control [allow | disallow | terminate] log [no | yes] }
    other-action { control [allow | disallow | terminate] log [no | yes] }
  }
}
description [string]
identity-users [add | delete | replace-all-with] {
  [identity user list names...]
}
identity-groups [add | delete | replace-all-with] {
  [identity group list names...]
}
}
}
rules none
actions [add | delete | modify] {
  [ [name] ] {
    options:
    shell-action { control [allow | disallow | terminate] log [no | yes] }
    sub-system-action { control [allow | disallow | terminate] log [no | yes] }
    sftp-up-action { control [allow | disallow | terminate] log [no | yes] }
    sftp-down-action { control [allow | disallow | terminate] log [no | yes] }
    scp-up-action { control [allow | disallow | terminate] log [no | yes] }
    scp-down-action { control [allow | disallow | terminate] log [no | yes] }
    rexec-action { control [allow | disallow | terminate] log [no | yes] }
    local-forward-action { control [allow | disallow | terminate] log [no | yes] }
    remote-forward-action { control [allow | disallow | terminate] log [no | yes] }
    x11-forward-action { control [allow | disallow | terminate] log [no | yes] }
    agent-action { control [allow | disallow | terminate] log [no | yes] }
    other-action { control [allow | disallow | terminate] log [no | yes] }
  }
}
auth-info [add | delete | modify] {
  [ [name] ] {
    options:
  }
  proxy-server-auth {
    private-key [string]
    public-key [string]
  }
}
```

```

}
proxy-client-auth {
  private-key [string]
  public-key [string]
}
real-server-auth {
  public-key [string]
}
}
}
timeout [integer]
lang-env-tolerance [any | common | default-value | none]

```

```

edit profile
options:
  all-properties
  non-default-properties

```

```

DISPLAY
list profile
show running-config profile
options:
  all-properties
  non-default-properties
  one-line

```

DESCRIPTION

You can use the profile component to configure a shareable and reusable set of ssh profile rules.

EXAMPLES

```

create profile profile1 auth-info add {
  auth1 {
    proxy-server-auth {
      private-key "abcd"
      public-key "1234"
    }
    proxy-client-auth {
      private-key "efgh"
      public-key "5678"
    }
  }
}

```

Creates a ssh profile with auth-info of client facing auth and server facing auth.

```

modify profile profile1 actions add {
  action1 {
    sftp-up-action {
      control disallow log yes
    }
    shell-action {
      control terminate log yes
    }
  }
}

```

Modify existing profile by adding default actions of shell action and sftp action.

```

modify profile profile1 rules add {
  rule1 {
    actions add {
      action1 {
        sftp-up-action {
          control disallow log yes
        }
        shell-action {
          control terminate log yes
        }
      }
    }
    identity-groups add {
      "grp1" "grp2"
    }
    identity-users add {
      "usr1" "usr2"
    }
    description "rule1 and action1"
  }
}

```

Modify existing profile by adding rule1 to it with command actions and user and group identity info.

```
list profile
```

Displays the current list of profile rules.

OPTIONS

```

description
User defined profile description.

```

timeout

User defined timeout value.

lang-env-tolerance

Set the tolerance level for LANG environment variable settings. Only applicable when "other-action" is set to "disconnect" or "terminate". "common" allows only "en_US.UTF-8", while "any" allows any standard locale.

rules

Adds, deletes, or replaces a profile rule.

description

User defined rule description.

actions

Specifies the rule actions that the system takes when a profile is applied.

shell-action

Specifies the rule shell action info.

sub-system-action

Specifies the rule sub system info.

sftp-up-action

Specifies the rule sftp up action info.

sftp-down-action

Specifies the rule sftp up action info.

scp-up-action

Specifies the rule scp up action info.

scp-down-action

Specifies the rule scp up action info.

rexec-action

Specifies the rule rexec action info.

local-forward-action

Specifies the rule local forward action info.

remote-forward-action

Specifies the rule local forward action info.

x11-forward-action

Specifies the rule x11 forward action info.

agent-action

Specifies the rule agent action info.

other-action

Specifies the rule other action info.

identity-users

Specifies the rule users identity.

identity-groups

Specifies the rule groups identity.

actions

Specifies the profile default actions that the system takes when a profile is applied.

shell-action

Specifies the rule shell action info.

sub-system-action

Specifies the rule sub system info.

sftp-up-action

Specifies the rule sftp up action info.

sftp-down-action

Specifies the rule sftp up action info.

scp-up-action

Specifies the rule scp up action info.

scp-down-action

Specifies the rule scp up action info.

rexec-action

Specifies the rule rexec action info.

local-forward-action

Specifies the rule local forward action info.

remote-forward-action

Specifies the rule local forward action info.

x11-forward-action
Specifies the rule x11 forward action info.

agent-action
Specifies the rule agent action info.

other-action
Specifies the rule other action info.

auth-info
Specifies the authentication info of public key and private key for this profile.

proxy-server-auth
Specifies a set of private/public keys that can be used to authenticate proxy (BigIP) host server to the real clients during the initial key exchange of the SSH session between real clients and BigIP acting as a proxy server. A SSH Profile MUST have at least 1 set of private/public key configured for proxy server authentication.

private-key
Specifies the private key of the authentication algorithm (RSA, DSS etc) used for the proxy server authentication.

public-key
Specifies the public key of the authentication algorithm (RSA, DSS etc) used for the proxy server authentication.

proxy-client-auth
Specifies a set of private/public keys that can be used to support 'publicKey' based client authentication as defined in RFC 4252 (The Secure Shell (SSH) Authentication Protocol). Note that this is optional in a SSH profile and is only required to support 'publicKey' based client authentication (as defined in section 7 of the above mentioned RFC).

private-key
Specifies the private key of the authentication algorithm (RSA, DSS etc) used for the proxy client authentication.

public-key
Specifies the public key of the authentication algorithm (RSA, DSS etc) used for the proxy client authentication.

real-server-auth
Specifies public key that can be used to authenticate real host server to the proxy (BigIP) client during the initial key exchange of the SSH session between BigIP acting as a proxy client and a real ssh (backend) server. If user does not configure any public key for the real server authentication in a SSH profile, all (backend) real servers are always trusted.

public-key
Specifies the public key of the authentication algorithm (RSA, DSS etc) used for the real server authentication.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2016. All rights reserved.

BIG-IP 2017-10-30 security ssh profile(1)

security zone

NAME
zone - Configures firewall zones.

MODULE
security

SYNTAX
Zones are reusable objects that are used to classify traffic in firewall policy. Zone is defined as consisting of one or more Vlans, and traffic matching one of the zone member Vlans belongs to the Zone. Zone object can be used as a "source" or "destination" specifier in Firewall policy rules to either mean originating from, or destined-to traffic. Modify the zone component within the security zone module using the syntax shown in the following sections.

CREATE/MODIFY
create zone [name]
options:

```
copy-from [string]
modify zone [name]
options:
  vlans [add | delete | modify | replace-all-with] {
    [ vlan_name ]
  }
  vlans none
```

edit zone

```
DISPLAY
list zone
show running-config zone
options:
  all-properties
  non-default-properties
  one-line
```

DESCRIPTION

You can use the zone component to configure a shareable and reusable set of network firewall zones which can be associated as enforced or staged with a number of configuration objects of the following types: security firewall policy.

EXAMPLES

```
modify zone vlans add {
vlan-1 { }
vlan-2 { } }
```

Creates a zone configuration that includes vlan-1 and vlan-2 as members.

```
list zone
```

Displays the current list of zones.

OPTIONS

`copy-from`
(CREATE) Specifies the name of an existing policy from which to copy all configuration options.

`vlans`
Adds, deletes, or replaces a zone vlan member. Specifies one or more vlans against which the packet will be compared, when used with security firewall policy rules.

SEE ALSO

create, edit, list, modify, security firewall policy, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-06-27 security zone(1)

sys

sys air-filter-reset

NAME

air-filter-reset - Resets the air filter alert timer

MODULE

sys

SYNTAX

Reset the air filter alert timer air-filter-reset component within the sys module using the following syntax.

DISPLAY

```
run air-filter-reset
```

DESCRIPTION

You can use the air-filter-reset component to reset the air filter alert time. The primary usage of this command is to reset the timer after the air filter has been replaced.

EXAMPLES

```
run air-filter-reset
```

Resets the air filter alert timer.

SEE ALSO
tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2017-03-07 sys air-filter-reset(1)

sys alert lcd

NAME

alert - Manages outstanding alerts.

MODULE

sys

SYNTAX

Manages the outstanding alerts on the system when used with the syntax below.

DISPLAY

show alert lcd

MODIFY

reset-stats alert lcd

reset-stats alert lcd [priority { alert | critical | emergency | error | info | warning }]

DESCRIPTION

You can use the lcd component to display or clear the lcd alerts for the system. The clear lcd alerts can be issued with the priority option followed by the priority value in which case only alerts with matching priority are cleared.

EXAMPLES

show alert lcd

Displays lcd alerts.

reset-stats alert lcd

Clears all lcd alerts.

SEE ALSO

reset-stats, show, sys alert lcd, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013-2015. All rights reserved.

BIG-IP 2016-03-14 sys alert lcd(1)

sys aom

NAME

aom - Manages Always On Management subsystem

MODULE

sys

SYNTAX

Manages the Always On Management(AOM) subsystem when used with the syntax below.

DISPLAY

aom list

MODIFY

aom [readonly | webui | media-redirection | vkvm | ipmi] [enabled | disabled]

DESCRIPTION

You can use the aom component to manage the AOM access for the system. All AOM controls are disabled by default. These commands are not supported on all platforms.

readonly reduces the AOM menu to status only, the AOM cannot be used to change the system.

webui controls access to the AOM web services, when enabled this opens port 443 on the AOM for HTTPS. This option is not available on all platforms.

media-redirection controls if the AOM can access a virtual media server to allow the user to mount a file or iso image from a remote system to the AOM. Only allowed when webui is enabled. Media-redirection opens up ports 5120/5124(CDMEDIA) 5122/5126(FDMEDIA) and 5123/5127(HDMEDIA) on the AOM. This option is not available on all platforms.

vkvm controls AOM webui Virtual Keyboard, Video and Mouse redirection. Only allowed when webui is enabled. This opens up port 7578/7582. This option is not available on all platforms.

ipmi controls ipmi access over LAN. When enabled ipmi commands into the AOM over the network port 623(IPMI) and 52123(SSHSOL) are allowed.

EXAMPLES

```
list sys aom all-properties
```

Displays aom configuration.

```
modify sys aom ipmi enabled
```

Allows IPMI access into the Always On Management subsystem.

SEE ALSO

modify, list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-03-29 sys aom(1)

sys appiq config

NAME

config - Configures this BIG-IP for AppIQ feature.

MODULE

sys appiq

SYNTAX

Automatically configures AppIQ feature on this box with given ECM host IP.

MODIFY

```
modify config
```

options:

```
host-ip [ip address]
```

DISPLAY

```
list config
```

options:

```
all-properties
```

DESCRIPTION

When run for the first time, this command automatically creates necessary configuration objects to establish communication between BIG-IP and ECM cluster. The command can be reused to change the configured host IP of the ECM cluster.

OPTIONS

host-ip

Specifies an IPv4/6 address for the ECM host. By default, this value is "any6" and AppIQ is disabled. Upon setting this value, the feature is enabled.

User can later delete automatically created and configured AppIQ logging profiles under System configuration to disable the feature. Please note that the command needs to be re-run to enable the feature if the configuration is manually removed/changed.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-11-30 sys appiq config(1)

sys application apl-script

NAME

apl-script - Provides scripts that can be included by an application template.

MODULE

sys application

SYNTAX

Configure the apl-script component within the sys application module using the syntax in the following sections.

EDIT

```
create apl-script [name]
modify apl-script [name]
options:
  apl-checksum [[string] | none]
  apl-signature [[string] | none]
  description [[string] | none]
  ignore-verification [true | false]
edit apl-script [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
```

DISPLAY

```
list apl-script
list apl-script [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete apl-script [name]
```

GENERATE

Note: generate cryptographic signature or checksum based on apl script text.

```
generate sys application apl-script [name]
options:
  checksum
  signature
```

DESCRIPTION

An APL script contains APL that can be directly included into application templates. APL scripts provide a convenient way to build libraries of common presentation elements. For detailed description of application presentation language elements, See help page of sys application template

EXAMPLES

The following is a fairly simple example of an APL script and a template that makes use of the APL script. The APL script defines a user type that can then be used multiple times in different templates.

```
sys application apl-script com.f5.apl.example {
  define string port validator "PortNumber"
}
```

```
sys application template example_template {
  actions {
  definition {
    presentation {
      include "com.f5.apl.example"
    section my_section {
      string address1
      port portnum1
      string address2
      port portnum2
    }
  }
}
```

```
generate my_script checksum
```

Generate a checksum for the script text and add the checksum as a property.

```
generate my_script signature signing-key my_key
```

Generate a signature for the script text using the specified private key and add the signature as a property.

Note: For a script which includes a checksum or signature to successfully load, the script text contents must match the stored checksum or signature.

To temporarily stop the verification of signature or checksum and still retain the checksum or signature, the ignore-verification attribute must be set to true. This is done by editing the script and adding the ignore-verification attribute.

To completely clear the signature or checksum, simply set the attribute script-signature or script-checksum to empty string "". By doing so, the script will be processed as if it was never signed or checksummed.

```
modify apl-script my_script {
description none
script {
}
ignore-verification true
script-checksum 74778e7b13016e0b9329a17f8d2da601
total-signing-status checksum
verification-status checksum-verified }
```

OPTIONS

You can use these options with the apl-script command:

checksum

Generate a checksum for the script text and add the checksum to the script as a property. Only for use with the generate command.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

script

Contains the APL text that can be imported into application templates.

signature

Generate a signature for the script text using the specified private key and add the signature to the script as a property. Only for use with the generate command.

signing-key

The private key to use for signing the script. Only for use with the signature option.

SEE ALSO

create, delete, edit, glob, list, modify, regex, sys application template and generate.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010, 2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 sys application apl-script(1)

sys application custom-stat

NAME

custom-stat - Provides derived statistics for iStats.

MODULE

sys application

SYNTAX

Configure the custom-stat component within the sys application module using the syntax in the following sections.

EDIT

create custom-stat [key]

modify custom-stat [key]

options:

app-service [[string] | none]

keyspace [string]

formula [string]

measure [string]

edit custom-stat [[key] | [glob] | [regex]] ...]

options:

all-properties

DISPLAY

list custom-stat

list custom-stat [[key] | [glob] | [regex]] ...]

DELETE

delete custom-stat [key]

DESCRIPTION

Statistics are derived for objects in the given keyspace based on the given formula, producing the given measure.

EXAMPLES

```
create sys application custom-stat myKey
```

```
keyspace sys.application.service
```

```
measure conns_per_min
```

```
formula "rate counter conns 60"
```

Creates a derived iStat.

OPTIONS

You can use these options with the custom-stat component:

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

formula

Specifies the first token in the formula indicates the computation to be made. Currently only rates are supported.

rate

rate computes the rate of change of the source measure over the last seconds.

This is applicable only to numeric measures. The derived measure is of type gauge.

keyspace

Specifies that a derived iStat will be computed for all objects in the given keyspace for which the formula is computable (the source measure of the correct type exists).

measure

Specifies the name of the derived measure to be created. The type of the derived measure is dependent on the formula.

SEE ALSO

create, modify, sys application service

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2012. All rights reserved.

BIG-IP 2012-10-19 sys application custom-stat(1)

sys application service

NAME

service - Configures traffic management application services.

MODULE

sys application

SYNTAX

Modify the service component within the sys application module using the syntax shown in the following sections.

```
CREATE/MODIFY
create service [name]
modify service [name]
options:
description [string]
device-group [[string] | default | non-default | none]
execute-action [name]
lists [add | delete | modify | replace-all-with] {
  [name] {
    options:
value { [string]... }
value none
encrypted [yes | no]
  }
}
lists none
strict-updates [disabled | enabled]
tables [add | delete | modify | replace-all-with] {
  [name] {
    options:
column-names { [name] ... }
encrypted-columns { [name] ... }
rows { { row { [value] ... } row { [value] ... } ... } }
rows none
  }
}
tables none
template [name]
traffic-group [[string] | default | non-default | none]
variables [add | delete | modify | replace-all-with] {
  [name] {
    options:
value [string]
encrypted [yes | no]
  }
}
variables none

metadata
[add | delete | modify] {
  [metadata_name ... ] {
value [ "value content" ]
persist [ true | false ]
  }
}
```

```
edit service [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties
```

```
DISPLAY
list service
list service [ [ [name] | [glob] | [regex] ] ... ]
show running-config service
show running-config service [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties
one-line
partition
```

Note: Application Service objects are always created in a subfolder of the current folder. Make sure the correct path is used to display the service. E.G list myapplication.app/myapplication

```
DELETE
delete service [name]
```

OPTIONS

You can use these options with the service component:

description
User defined description.

device-group
Specifies the name of the device group to which the application service is assigned. If this property is modified with the default keyword, the value of the parent folder or partition will be used and the inherited-devicegroup property will be set to true.

execute-action
Runs the specified template action associated with the service.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

lists

Provides the set of list variables and values that are passed to template scripts.

metadata

Associates user defined data, each of which has name and value pair and persistence. The default value is persistent, which means the data will be saved into the config file.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

strict-updates

Specifies whether configuration objects contained in the application service can be directly modified outside the context of the system's application service management interfaces.

tables

Provides the set of table variables and values that are passed to template scripts.

template

The template defines the configuration for the application service. Generic application service has no template associated with it. This can be changed after the service has been created to move the service to a new template. A templated application service can be converted to a generic application service by setting new template to none or empty string. Similarly a generic application service can be made templated by associating it with the existing template.

template-modified

Indicates that the application template used to deploy the service has been modified. The application service should be updated to make use of the latest changes.

template-prerequisite-errors

Indicates any missing prerequisites associated with the template that defines this application.

traffic-group

The name of the traffic group that the application service is assigned to. If this property is modified with the "default" keyword, the value of the parent folder or partition will be used and the inherited-trafficgroup property will be set to true.

variables

The set of atomic variables and values that are passed to template scripts.

SEE ALSO

create, delete, edit, glob, list, modify, regex, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2014-07-28 sys application service(1)

sys application template

NAME

template - Enables the creation of user-defined templates.

MODULE

sys application

SYNTAX

Configure the template component within the sys application module using the syntax in the following sections.

CREATE/EDIT/MODIFY

create template [name]

modify template [name]

options:

actions [add | delete | modify | replace-all-with] {

definition {

options:

html-help [string]

implementation { [string] }

presentation { [string] }

role-acl [add | delete | modify | replace-all-with] {

[role]

}

role-acl none

```

run-as [string]
  }
}
description [string]
requires-modules [add | delete | modify | replace-all-with] {
  [string]
}
requires-bigip-version-max [string]
requires-bigip-version-min [string]
metadata
  [add | delete | modify] {
    [metadata_name ... ] {
value [ "value content" ]
persist [ true | false ]
    }
  }
}

```

edit template [name]

DISPLAY
list template
list template [name]

DELETE
delete template [name]

GENERATE

Note: generate cryptographic signature or checksum based on template fields - html-help, implementation, macro and presentation in definition section.

generate template [name]

options:
checksum
signature

SAVE/LOAD
save template [name] file [filename]
load template [filename]

DESCRIPTION

Application templates allow a user to define a custom interface for easily creating complex configurations. The user can create multiple templates for various types of configurations. Once built, the user can use a template to create an application, which is a specific set of configuration objects (such as Virtual IP addresses, pools, and so forth), that work together to perform some task.

The template is composed of two primary parts, the presentation and the implementation.

The presentation section describes a form (a set of questions and user interface elements) that the user must fill out in order to create an application.

The implementation section describes how the values collected from the user (the form variables) are used to generate the actual configuration objects which are part of the application.

The presentation section of the template is written in a simple language called Application Presentation Language or APL. The implementation section of the template is written in TCL and provides access to tmsh scripting commands.

APPLICATION LIFECYCLE

Before describing in detail how a template is written, it is important to explain how the resulting template will be used. Since templates are used to create and edit applications, it makes sense to review the application lifecycle.

Application Creation

The user selects which template to use for his application. The system presents an empty form, based on the template's presentation script that the user fills out and submits. The system collects and stores the form variables in a newly created application object. Configuration objects are generated based on the form variables by the template's implementation script.

Application Editing

The user selects an existing application that he would like to change. The system reloads the form associated with the template that was used to create the application and refills all form variables using the previous user input, which is gathered from the application object. The user edits the form and submits it. The template's implementation script is run again to compute a new set of configuration objects for the application. The system alters the current configuration objects associated with the application to match the newly computed set of configuration objects, including creating, modifying, and deleting objects as needed.

Application Deletion

The user selects an application to delete. All configuration objects associated with the application are removed.

APPLICATION TEMPLATE LANGUAGE

The application template language describes the user interface presented to a user making a new application, or editing an existing application. It describes what questions to ask, how the questions are presented (for example, a free form field or a list of options), and the names of the variables used to store the values the user inputs.

It consists of a set of primitive form elements (string, choice, etc), a set of grouping and organization

constructs (section, table, etc), methods for hiding or displaying portions of the form based on the values of other portions (optional), a method to associate human-readable text with various form elements (text) and methods for creating user defined types(define group, define section, etc) for reuse of application presentation language elements.

Primitive Elements

Primitive elements represent the actual user interface components. The system displays each primitive element as part of the form, and associates it with a form variable. The following lists the basic primitive types:

choice

A list of options from which the user can select (a drop-down menu).

```
choice [default "" ] [display "" ] { "", "", ... }
```

editchoice

Multiple choices are available that the user can select, or a new value can be entered if the choices are not acceptable.

```
choice [default "" ] [display "" ] { "", "", ... }
```

multichoice

Similar to a basic choice element except that multiple items may be selected from the available choices.

```
choice [default "" ] [display "" ] { "", "", ... }
```

password

Similar to a string element except the contents may be obscured to prevent others from seeing the value.

```
password [default "" ] [display "" ] [required]
```

string

A basic text box into which the user can enter an arbitrary string.

```
string [default "" ] [display "" ] [required] [validator "" ]
```

Each primitive element is associated with a variable name, which defines where the value collected by the form is stored. In addition, primitive elements can have additional parameters such as a default value, a validation method that provides for additional requirements (for example, the string must be an IP address).

The following defines the format for the string primitive values, using normal BNF syntax:

Â· default - A sensible default value to which the string is initialized when a new application is created.

Â· display - Directs the renderer how to display the element. This can be small, medium, large, xlarge, or xxlarge.

Â· required - If present, a valid value must be entered before the application can be created.

Â· validator - The name of a well known validation method.

Section

The section construct is used to group form variables (primitives) into logical sections for display.

Each section is named, and header text can be defined for a section using the text construct.

Every variable must be inside a section. The format for a section is:

```
section { }
```

For example, to represent the data associated with a virtual IP:

```
section vip
{
  string address
  string port default "80" display "small"
}
```

Table

The table construct is similar to section, except that it represents a grouping of elements that can be repeated zero or more times. The syntax for table and section are identical.

```
table { }
```

For example, to collect a list of nodes from a user to populate a pool, you can add any number of nodes, each of which has an address and port:

```
section pool
{
  table members
  {
    string address
    string port default "80" display "small"
  }
}
```

The table above is displayed using a JavaScript-editing widget that enables you to add and remove pool members. Each member contains two form variables: address and port.

Optional

The optional construct allows the form elements to be hidden or shown based on the state of other form elements. The syntax of the optional construct is:

```
optional () { }
```

The expression in the optional construct is evaluated during the display of the form. The content section is displayed or hidden, based on its value.

To dynamically hide parts of the presentation based on the answer to a earlier question, use the variable name in the expression:

```
section chooseopts {
  choice show_section_1 {"yes", "no"}
}
section section1
{
  optional (chooseopts.show_section_1 == "yes")
  {
    string str
  }
}
```

User defined types

The define construct allows the creation of user-defined types out of primitive types. The defined type can then be used multiple times independently at different places. This is especially useful in conjunction with the include element because types can be defined in the included application presentation language script and then used where necessary in the template. For more details on application presentation language script, See help sys application apl-script.

For example, user defined choice type can be defined as below and can be reused at multiple sections:

```
define choice yesno {
  "Yes", "No"
}
section ssl_section {
  yesno use_ssl
}
section optimizations {
  yesno use_wa
  yesno offload_ssl
}
```

The define group construct allows the creation of user-defined type to allow the user to group multiple elements of existing types together. The defined type can be reused multiple times independently similar to the above.

For example IpAddress and port can be grouped into a user-defined type and reused in multiple sections:

```
define group addrport {
  string addr required validator "IpAddress"
  string port
}
section http_section {
  addrport server
}
section sip_section {
  addrport client
  addrport server
}
```

Localization

The text element lets you define the localized text labels for sections, table, row and other sub-elements.

For message element, body text can be localized in addition to the label. Similarly for the choice, editchoice and multichoice element, display text associated with each choice value can be localized. The syntax for the text element is:

```
text [""] {
  ""
  . ""
  . "" ""
  . "" { "" => "", "" => "", ... }
}
```

Depending on the locale used (setting in the browser), particular text label, body text or choice display text will be shown to the user.

For example, string, message and choice display texts can be localized as below.

```
section http
{
  message intro
```

```

string address
string port default "80" display "small"
choice pools default "pool1" { "pool1", "pool2", "pool3" }
choice profile default "http" tcl {
    set choices "no\n"
    append choices "http\n"
    append choices [tmsh::run_proc f5.app_utils:get_items ltm profile http]
    return $choices
}
}

text {
vs "HTTP Application"
vs.intro "Introduction" "This template supports simple web server implementations"
vs.address "What IP address do you want to use for this virtual server?"
vs.port "What port do you want to use for this virtual server?"
vs.pools "Use pool.." {"Internal" => "pool1", "Public cloud" => "pool2", "Private data center" => "pool3" }
vs.profile "Use profile.." { "Do not use profile" => "no", "Use F5's recommended profile" => "http" }
}

text "de_AU" {
vs "HTTP-Anwendung"
vs.intro "Einführung" "Diese vorlage unterstützt einfache web-server-implementierungen"
vs.address "Welche ip-adresse mochten Sie für diesen virtuellen Server zu verwenden?"
vs.port "Welchen port willst du für diesen virtuellen Server zu verwenden?"
vs.pools "Verwenden pool.." {"intern" => "pool1", "Privat Rechenzentrum" => "pool3", "Öffentliche Cloud" => "pool2" }
vs.profile "Mit profil.." { "Verwenden sie kein profil" => "no", "Verwenden von F5 empfohlen profil" => "http" }
}

```

A user from Austria will see the german text, all other locales will see the default (locale-less) text.

While localizing choice value display text, users are allowed to use different ordering of choice values in each locale. If TCL is used to populate the choices, then best effort is made to match what is returned in the TCL to the given localized choice value. In the above example, the embedded TCL script for profile will return two static choices (no and http) followed by the list of all http profiles. These static choices are localized, but not the other results. When the TCL results contain a mix of localized and non-localized choices, the localized choices will always be listed first in the order specified in the text element.

With the localization, message body and static choices will become optional in the declaration. If the message body is provided in both the declaration and in text element, the body in the text element will override the body in the variable declaration. Same applicable for the display text of choice value provided in declaration. The recommended syntax for choice, editchoice and multichoice element is to give just the choice values in the variable declaration, and give the display text of the choices in the text element.

TMSH SCRIPTING SUPPORT

Once the user finishes editing an application, the form variables are saved, and the implementation section of the associated template is run. The implementation section is an ordinary TCL script and can use the standard set of tmsh scripting extensions. In addition, there are a few template-specific additions.

First, access to the form variable is done using the syntax, where `is` is the name of the section to which the variable belongs, and `is` is the name of the form variable:

```
$::
```

Next, a table can be iterated over, and for each list element, the components of the list can be gathered using the `tmsh::get_field_value` command. For example, for the pool member example described in the section regarding the list, you can use the following syntax:

```

foreach member $::pool__members {
set the_addr [tmsh::get_field_value $member address]
set the_port [tmsh::get_field_value $member port]
# Do something with the_addr and the_port
}

```

This means for variable access can also be used within a script macro. Expansion of a macro is done using the `tmsh::expand_macro` command. Usage:

```
tmsh::expand_macro [macro] [name_value_pair_list]
```

The variables defined in the presentation are automatically available from within the macro. If additional variables are needed from within the macro, they can be specified via `name_value_pair_list`. Variables defined this way will take precedence over duplicate variables defined in the presentation.

TMSH BUILT-IN VARIABLES

Specific details on application and application template is provided to implementation section using built-in variables. Following are the variables available for use.

`tmsh::app_name`
Stores the user-provided application name string.

`tmsh::app_name_path`
Stores the path name of application in configuration database.

`tmsh::app_template_name`
Stores the user-provided application template name including the path in configuration database.

tmsh::app_template_action
Stores the application template action name.

EXAMPLES

The following template example shows both the presentation and implementation sections. (It lacks some features, such as use of optional, defaults, validators, etc.)

```
presentation {
  section basic
  {
    choice ssl_enabled { "true", "false" }
    string addr
    string more_stuff
    table servers
    {
      string addr
      string port
      string ratio
    }
  }
}
text
{
  basic "Some example questions"
  basic.ssl_enabled "Should SSL be enabled?"
  basic.addr "What address should we use for the VIP?"
  basic.servers.addr "Address"
  basic.servers.port "Port"
}
}

implementation {
  if { $::basic_ssl_enabled }
  {
    set profile_name [format "%s_%s" $tmsh::app_name clientssl]
    tmsh::create ltm profile client-ssl $profile_name
    append profile_name " http"

    set destination "$::basic_addr:https"
    set monitor https
  }
  else
  {
    set profile_name http
    set destination "$::basic_addr:http"
    set monitor http
  }
}

set pool_name [format "%s_%s" $tmsh::app_name pool]

set members \{
  foreach server $::basic_servers {
    append members [tmsh::get_field_value $server addr]
    append members ":"
    append members [tmsh::get_field_value $server port]
    append members " { ratio "
    append members [tmsh::get_field_value $server ratio]
    append members "}"
    append members " "
  }
}
append members \}

tmsh::create ltm pool $pool_name \
members replace-all-with $members \
monitor $monitor

set vs_name [format "%s_%s" $tmsh::app_name virtual]
tmsh::create ltm virtual $vs_name \
destination $destination \
profiles replace-all-with "{ $profile_name }" \
snat automap \
pool $pool_name \
http-class none
}
```

generate my_app checksum

Generate a checksum for the template definition and add the checksum as a property.

generate my_app signature signing-key my_key

Generate a signature for the template definition using the specified private key and add the signature as a property.

Note: For a template which includes a checksum or signature to successfully load, the definition contents must match the stored checksum or signature.

To temporarily stop the verification of signature or checksum and still retain the checksum or signature, the ignore-verification attribute must be set to true. This is done by editing the script and adding the ignore-verification attribute.

To completely clear the signature or checksum, simply set the attribute script-signature or script-checksum to empty string "". By doing so, the script will be processed as if it was never signed or checksummed.

```
sys application template my_tmpl {
actions {
  definition {
html-help {

}
implementation {
  # insert tmsch script
}
presentation {
  # insert apl script
}
role-acl none
run-as none
}
description "This is my template"
ignore-verification true
script-checksum 74778e7b13016e0b9329a17f8d2da601
total-signing-status checksum
verification-status checksum-verified }
```

OPTIONS

actions

Adds, deletes, or replaces a set of template actions. You can configure the following options for an action:

html-help

The help for the application template action formatted as HTML.

implementation

The script that is run to create the configuration objects associated with the application.

name The name of the application template action.

presentation

The questions that must be answered to create an application from the template.

role-acl

The list of roles that are allowed to run the action.

run-as

The user account that will be used to run the implementation script. If no account is specified, the script is run as the calling user.

checksum

Generate a checksum for the template definition and add the checksum to the template as a property. Only for use with the generate command.

signature

Generate a signature for the template definition using the specified private key and add the signature to the template as a property. Only for use with the generate command.

signing-key

The private key to use for signing the template. Only for use with the signature option.

description

User defined description.

metadata

Associates user defined data, each of which has name and value pair and persistence. The default value is persistent, which saves the data into the config file.

partition

Displays the administrative partition within which the application template resides.

prerequisite-errors

A message indicating if there are any errors with the prerequisites for the template on the current BIG-IP system. If there are errors no applications can be created from this template. If there are no errors then the template is valid.

requires-modules

Adds, deletes, or replaces the list of modules that are required to be provisioned for this template to work.

requires-bigip-version-max

Specifies the maximum version of BIG-IP software required by this template.

requires-bigip-version-min

Specifies the minimum version of BIG-IP software required by this template.

THIRD PARTY TCL LIBRARY USAGE

A selection of third party libraries have been tested to work within the CLI script environment. These include MD5, BASE64, SHA1/SHA256, HTTP, TLS, TCL Perl, LDAP client, and XML parser. The TCL packages can only reside in the /use/share/compat-tcl8.4 directory.

Important: Only these tested packages are supported currently.

The following example shows how the Tcl package command can make use of the XML parser:

```
cli script /Common/use_xml {

proc script::EStart {tag attlist args} {
  array set attr $attlist
  puts "Element \"$tag\" started with [array size attr] attributes"
}

proc script::PCData text {
  incr ::count [string length $text]
}

proc script::run {} {
  namespace eval :: {
    set count 0
  }
  puts "running use_xml...\n"
  set pkg_name xml
  if {[catch {package require $pkg_name 3.2}]} {
    puts "No package found: $pkg_name!\n"
  }
  else {
    puts "Found package: $pkg_name!\n"
    set parser [xml::parser]
    $parser configure -elementstartcommand script::EStart -characterdatacommand script::PCData
    set fp [open "/shared/test.xml" r]
    set text [read $fp]
    $parser parse $text
    puts "The document contains $::count characters"
    close $fp
  }
}
```

Here are some additional examples:

```
cli script /Common/use_sha1 {
proc script::run {} {
  set pkg_name sha1
  if {[catch {package require $pkg_name}]} {
    puts "No package found: $pkg_name!\n"
  }
  else {
    puts "Found package: $pkg_name!\n"
    puts "TCL does SHA1 now:"
    puts [sha1::sha1 -hex "TCL does SHA1"]
  }
}
```

```
cli script /Common/use_base64 {
proc script::run {} {
  set pkg_name base64
  if {[catch {package require $pkg_name}]} {
    puts "No package found: $pkg_name!\n"
  }
  else {
    puts "Found package: $pkg_name!\n"
    set chemical [encoding convertto utf-8 "C\u2088H\u2081\u2080N\u2084O\u2082"]
    set encoded [base64::encode $chemical]
    set caffeine [encoding convertfrom utf-8 [base64::decode $encoded]]
    puts "Caffeine: $caffeine"
  }
}
```

SEE ALSO

edit, list, modify, show, tmsh, generate

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015-2016. All rights reserved.

sys autoscale-group

NAME
autoscale-group - Configure autoscale group ID for BIG-IP VE Autoscale Service on Amazon Web Services(AWS).

MODULE
sys

SYNTAX
Configure the autoscale-group component within the sys module using the syntax in the following sections.

CREATE/MODIFY

modify sys autoscale-group

Properties:

autoscale-group-id [[string] | none]
description [string]

edit sys autoscale-group

Options:

all-properties
non-default-properties

DISPLAY

list sys autoscale-group

Options:

all-properties
non-default-properties
one-line

Properties:

autoscale-group-id
description

DESCRIPTION

Specifies Amazon Web Services (AWS) Auto-Scaling Group ID to which given BIG-IP-VE belongs to. These settings will be used for Auto Scaling BIG-IP instances based on user specified policy by Amazon Web Services(AWS).

OPTIONS

autoscale-group-id

Specifies autoscale-group id as reported by Amazon Web Services(AWS).

description

User defined description.

SEE ALSO

edit, list, modify, sys autoscale-group, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2017-08-15 sys autoscale-group(1)

sys availability

NAME
availability - Display summary of system availability.

MODULE
sys availability

SYNTAX
Display summary of system availability.

DISPLAY

show sys availability

options:
field-fmt

RESET

reset-stats sys availability

DESCRIPTION

Displays a summary of the availability metrics for the BIG-IP(r) system or blade. The metrics represent the

amount of time spent in a given state.

States

Available

The total time BIG-IP(r) has been active or standby for one or more traffic-groups.

Maintenance

The total time BIG-IP(r) has spent offline for maintenance, e.g., state administratively set to offline, rebooting, powered off. This will also include the time running in software which does not log system availability.

Unavailable

The total time BIG-IP(r) is operationally offline and unable to accept traffic or failover.

Unknown

The total time for which the system is unable to determine availability state. This could be due to gaps present in availability logs or when the system is unable to write to the availability log.

EXAMPLES

```
show sys availability
```

Displays summary of system availability metrics.

```
reset-stats sys availability
```

Resets system availability metrics.

OPTIONS

```
field-fmt
```

Displays the summary of system availability metrics in field format where the value is total time in seconds spent in given state.

SEE ALSO

sys failover, cm traffic-group, reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017-2018. All rights reserved.

BIG-IP 2018-10-20 sys availability(1)

sys clock

NAME

clock - Displays the current date and time.

MODULE

sys

SYNTAX

DISPLAY

```
show clock
```

options:

```
field-fmt
```

```
modify clock
```

options:

```
time [time]
```

DESCRIPTION

You can use the clock component to display the system date and time.

EXAMPLES

```
show clock
```

Display the current date and time.

```
modify clock time 2012-12-11:12:30:45
```

Set the system clock to the specified time.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009, 2012-2013. All rights reserved.

BIG-IP 2013-03-22 sys clock(1)

sys cluster

NAME

cluster - Configures a cluster in a VIPRION(r) system.

MODULE

sys

SYNTAX

Configure the cluster component within the sys module using the syntax in the following sections.

MODIFY

modify cluster [name]

options:

address [IP address/Netmask | none]

alt-address [IP address/Netmask | none]

members {

[1 | 2 | 3 | 4 | 5 | 6 | 7 | 8] {

options:

address [IP address | none]

alt-address [IP address | none]

[disabled | enabled]

priming [disabled | enabled]

}

}

min-up-members [integer]

min-up-members-enabled [no | yes]

edit cluster default

options:

all-properties

non-default-properties

DISPLAY

list cluster

options:

all-properties

non-default-properties

one-line

show running-config cluster

show running-config cluster [option name]

options:

one-line

show cluster

show cluster [option name]

options:

all-properties

field-fmt

DESCRIPTION

You can use the cluster component to modify the configuration of the primary blade in a cluster. When you do this, the system automatically propagates the changes to the other blades in the cluster. This is known as cluster synchronization.

EXAMPLES

modify cluster default address 192.168.217.44/24

Sets the floating management IP address for the cluster default to an IP address of 192.168.217.44.

list cluster my_cluster

Displays the properties of the cluster named my_cluster.

OPTIONS

address

Specifies an IP address for the cluster or cluster member. The default value is none.

alt-address Specifies an optional, additional IP address for the cluster or cluster member in support of dual-stack addressing. The default value is none. Note: The IP address specified must compliment the family of the IP address held in address. I.e. Without loss of generality, if address holds an IPv6 address, then alt-address must hold an IPv4 address or none at all. Symmetrically, if address holds an IPv4 address, then alt-address must hold an IPv6 address or none at all.

disabled

Disables the specified cluster member. The default value is enabled.

enabled

Enables the specified cluster member. This is the default value.

members

Specifies the cluster members to be acted on by the command. A cluster member is a slot into which you insert a blade. The cluster member is identified by the number assigned to the slot.

min-up-members

Specifies the minimum number of cluster members that must be up for the cluster to remain Active. The default value is 0.

min-up-members-enabled

When set to yes, specifies that when the number of cluster members that are active is below the value of the option min-upmembers, the cluster fails over to its peer. The default value is no.

Enable this parameter when you configure a redundant pair.

Important: Make sure that you modify the value of the min-up-members option appropriately when you take blades down in a cluster. Otherwise, you can get into the condition where disabling a cluster member brings the cluster below the value of the option min-up-members, which can cause the cluster to fail over to its peer.

name

Specifies a name for the cluster. This option defaults to the value default.

priming

Prevents a cluster member from proceeding to the RUNNING cluster quorum state, which is useful when a blade is in a reboot loop. The default value is disabled.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2017-03-01 sys cluster(1)

sys config-diff

NAME

config-diff - Displays the differences between two specified single configuration files (SCFs).

MODULE

sys

SYNTAX

Display information using the config-diff component within the sys module with the syntax in the following section.

DISPLAY

show config-diff [file name] [file name]

DESCRIPTION

You can use the config-diff component to display the differences between two previously created SCF files.

EXAMPLES

show config-diff my.scf your.scf

Displays information about the differences between two specified files.

OPTIONS

file name

Specifies the name of an SCF file that you want to compare to another SCF file.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-11 sys config-diff(1)

sys config

NAME

config - Manages the BIG-IP(r) system configuration.

MODULE

sys

SYNTAX

Save the running configuration or load the system configuration files within the sys module using the following syntax.

MODIFY

save config

options:

base

binary

current-partition

exclude-gtm

file

gtm-only

no-passphrase

one-line

passphrase

partitions

tar-file

time-stamp

user-only

wait

load config

options:

base

current-partition

default

exclude-gtm

file

files-folder

from-terminal

gtm-only

merge

replace

passphrase

partitions

tar-file

user-only

verify

delete config file [file name]

DISPLAY

list config file

DESCRIPTION

The system applies all configuration changes that you make from within tmsh to the running configuration. To save the running configuration to the system configuration files, use the command sequence save config. Additionally, you can replace the running configuration with the configuration in the system configuration files using the command sequence load config.

If any of these options are not specified, save/load config will save or load the configuration in all partitions on this system:

- binary
- default
- file
- from-terminal
- partitions

EXAMPLES

save config

Saves the running configuration in all partitions by overwriting the system configuration files.

In Virtual Editions with f5-swap-eth installed, saves the mapping of Ethernet device names and MAC addresses to /etc/ethmap to make the working BIG-IP still work after adding/deleting virtual NIC(s). It also works for save config partitions all.

save config base

Saves the running base configuration in all partitions by overwriting the system base configuration files.

save config binary

Saves all running configuration by overwriting the system binary configuration database file.

save config current-partition

Saves the running configuration in current update partition by overwriting the system configuration files.

save config wait

Save request waits if another save operation is in progress.

save config file my_file tar-file my_tar_file no-passphrase

Saves all running configuration to the specified file, my_file, and all the user provided disk files referred to by the configuration into my_tar_file. From v12.1.0 onwards the no-passphrase option should be explicitly specified for saving the unencrypted SCF and SCF.tar.

save config file my_file passphrase my_password

Saves all running configuration to the specified file, my_file and encrypt it with my_password.

save config partitions { my_partition }

Saves the running configuration in my_partition by overwriting the system configuration files.

save config partitions all

Saves the running configuration in all partitions by overwriting the system configuration files.

save config user-only

Saves only user account configuration by overwriting the system configuration files.

load config

Replaces the running configuration in all partitions with the configuration in the system configuration files.

load config base

Replaces the running base configuration in all partitions with the configuration in the system base configuration files.

load config current-partition

Replaces the running configuration in current update partition with the configuration in the system configuration files.

load config merge file my_file

Loads the specified configuration from my_file, which modifies the running configuration.

load config replace file my_file

Loads the specified configuration from my_file, which modifies the running configuration, overwrites certain elements that may be conflicting.

load config verify file my_file

Validates the specified configuration in my_file to see whether they are valid to replace the running configuration. The running configuration will not be changed.

load config verify merge file my_file

Validates the specified configuration in my_file to see whether they are valid to be merged into the running configuration. The running configuration will not be changed.

```
load config verify replace file my_file
```

Validates the specified configuration in my_file to see whether they are valid to be merged into the running configuration with some elements replaced. The running configuration will not be changed.

```
load config default
```

Sets system configuration back to factory default settings.

```
load config file my_file tar-file my_tar_file
```

Replaces all running configurations with the configuration in the specified file, my_file and the disk files referred to by the configuration are retrieved from my_tar_file.

```
load config file my_file files-folder my_files_folder
```

Replace all running configuration with the configuration in the specified file, my_file and the disk files referred to by the configuration is taken from the directory tree under my_files_folder.

```
load config file my_file passphrase my_password
```

Replaces all running configuration with the configuration in the specified encrypted file, my_file and decrypt it with my_password.

While searching for disk files under the specified folder, the order of search is first by file name as in cache-path, and then by object-name. If more than one file is found for a name, then the relative path in the cache-path is used to make the selection.

That is, while looking for

Looks for file(s) named B.

If none are found, looks for file(s) named "xxx"

When more than one file is found, looks for a copy that matches paths in the order:

```
B>
```

```
B>
```

```
load config partitions { x }
```

Replace the running configuration in partition x with the configuration in the system configuration files.

```
load config partitions all
```

Replace the running configuration in all partitions with the configuration in the system configuration files.

```
load config from-terminal
```

Replace the running configuration with what is entered from the terminal.

1. Type the initial command. 2. The system responds with a confirmation prompt, type Y to confirm.

Replace the running configuration? (y/n) y

3. Type in the replacement configuration entries.

```
net self-allow {
  defaults {
    igmp:any
    ospf:any
    pim:any
    tcp:161
    tcp:22
    tcp:4353
    tcp:443
    tcp:53
    udp:1026
    udp:161
    udp:4353
    udp:520
    udp:53
  }
}
net stp-globals {
  config-name 00-01-D7-B5-67-00
}
sys management-ip 172.27.41.70/24 { }
sys management-route default {
  gateway 172.27.41.254
}
sys provision ltm {
  level nominal
}
```

```
....
ltm pool pool1 {
  slow-ramp-time 200
}
.....
^D
```

4. Use Ctrl+D to submit the changes or Ctrl+C to cancel the changes.

delete config file myfile

Delete myfile in default directory, /var/local/scf/.

list config file

Display files in default directory, /var/local/scf/.

OPTIONS

base Indicates the base configuration. This option cannot be specified with the binary, default, gtm-only, and user-only options.

binary
Indicates binary configuration. This option may not be specified with any other options.

default
Indicates factory default configuration. This option cannot be specified with any other options.

file Loads or saves a configuration from the specified file. For save, a file with a relative path is saved in the default directory, /var/local/scf. For load, in shell mode, the default directory, /var/local/scf, is used for a file with a relative path. In bash mode, for a file with a relative path, the current directory is searched first. If the file can't be found in the current directory, /var/local/scf is searched.

This option can be used with binary, default, from-terminal and partitions options.

passphrase
Specifies a password to save or load an encrypted configuration file. This option can only be used with option file.

tar-file
Loads or saves disk files referred to by the configuration from the specified tar file. A file with a relative path is looked up, relative to the current directory.

files-folder
Loads disk files referred to by the configuration from the folder tree under the specified folder. Disk files by name are searched for recursively. When there is more than one file with the same name, the relative path of the file from the cache-path is used for selection.

from-terminal
Specifies that the configuration will be input from the terminal in the same format as the system configuration files in . **Use Ctrl+D to submit the changes and Ctrl+C to cancel the changes.**

This option cannot be specified with default, file and partitions.

gtm-only
Indicates the Global Traffic Manage (GTM) configuration. This option cannot be specified with the base, exclude-gtm, and user-only options.

exclude-gtm
Indicates the BIG-IP configuration, excluding GTMs. This is only valid with the file option. This option cannot be specified with the base, gtm-only, and user-only options.

merge
Loads the configuration from the specified file or from the terminal, which modifies the running configuration. If merging from the terminal, it requires Ctrl+D to complete the operation. This option is only valid with the file or from-terminal options.

partitions
Indicates the partitions in which configuration components reside. This option cannot be specified with the default, file, from-terminal, or merge options.

user-only
Indicates the configuration including user account information only. This option cannot be specified with the base, default, exclude-gtm, or gtm-only options.

time-stamp
Inserts a time-stamp in a file name. This is only valid with the file option.

verify
Validates the specified configuration from file(s) or from the terminal without changing the running configuration.

wait Specifies that tmsh should wait for another instance of tmsh to finish saving the configuration before proceeding. If wait is not specified and another instance of tmsh is in the process of saving the configuration, the command exits tmsh immediately (because the other instance of tmsh is already saving the configuration).

SEE ALSO

load, save, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015-2016. All rights reserved.

BIG-IP 2019-01-14 sys config(1)

sys connection

NAME

connection - Sets idle timeout for, displays, and deletes active connections on the BIG-IP(r) system.

MODULE

sys

SYNTAX

Use the connection component within the sys module to manage connections using the following syntax.

MODIFY

modify connection

options:

idle-timeout [integer]
flow-accel-type software-only

DISPLAY

show connection

option:

all-properties
age [integer]
cs-client-addr [IP address/prefixlen]
cs-client-port [[integer] | [service]]
cs-server-addr [IP address/prefixlen]
cs-server-port [[integer] | [service]]
connection-id [integer]
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
protocol [name]
save-to-file [filename]
ss-client-addr
ss-client-port [[integer] | [service]]
ss-server-addr
ss-server-port [[integer] | [service]]
type [any | mirror | self | mptcp]

DELETE

delete connection

option:

age [integer]
cs-client-addr [IP address]
cs-client-port [[integer] | [service]]
cs-server-addr [IP address]
cs-server-port [[integer] | [service]]
connection-id [integer]
protocol [name]
ss-client-addr [IP address]
ss-client-port [[integer] | [service]]
ss-server-addr [IP address]
ss-server-port [[integer] | [service]]
type [any | mirror | self | mptcp]

DESCRIPTION

You can use the connection component to set the idle timeout for or delete active connections to the BIG-IP system based on a specified filter. Additionally, you can display information about the active connections to the system.

You can specify the option using either a number or a service (80 or http).

Important: If you do not specify a port or service, the system deletes all connections that match just the IP address. If you do not specify an IP address, the system deletes all connections including mirrored connections.

EXAMPLES

show connection all-properties

Displays information about all active connections to the system.

modify connection idle-timeout 300

Changes the amount of idle time before a connection is disconnected to five minutes (300 seconds).

modify connection flow-accel-type software-only

Flush all hardware accelerated connections from the hardware to be software only connections.

OPTIONS

age Specifies, in seconds, the minimum idle time of the active connections that you want to display or delete.

cs-client-addr/prefixlen

Specifies the client-side remote IP address of the active connections that you want to display or delete. Subnet addresses are allowed with prefix lengths up to /24 and /56 for IPv4 and IPv6 respectively.

cs-client-port

Specifies the clientside remote port of the active connections that you want to display or delete.

cs-server-addr/prefixlen

Specifies the clientside local IP address of the active connections that you want to display or delete. Subnet addresses are allowed with prefix lengths up to /24 and /56 for IPv4 and IPv6 respectively.

cs-server-port

Specifies the clientside local port of the active connections that you want to display or delete.

connection-id

Specifies the MPTCP connection ID of the active connections that you want to display or delete.

idle-timeout

Specifies the interval, in seconds, that a connection can remain idle before the system closes the connection.

protocol

Specifies the protocol of the active connections that you want to display or delete.

save-to-file

Specifies the file which connection information can be save to. With this option, it can write a file larger than 2GB.

ss-client-addr

Specifies the serverside local IP address of the active connections that you want to display or delete.

ss-client-port

Specifies the serverside local port of the active connections that you want to display or delete.

ss-server-addr

Specifies the serverside remote IP address of the active connections that you want to display or delete.

ss-server-port

Specifies the serverside remote port of the active connections that you want to display or delete.

type Specifies the type of active connections that you want to display or delete. The possible values are:

any Specifies all active connections.

mirror

Specifies only mirrored connections.

self Specifies the connection with which you are accessing the system.

mptcp

Specifies the MPTCP type of connections.

SEE ALSO

delete, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2015. All rights reserved.

BIG-IP 2018-05-31 sys connection(1)

NAME

console - Configures the serial console for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the console component within the sys module using the syntax in the following section.

MODIFY

modify console

options:

baud-rate [integer]

DISPLAY

show console

DESCRIPTION

You can use the console component to configure the serial console on the BIG-IP system.

OPTIONS

baud-rate

Specifies the baud rate for the serial console. Select from the following options:

• 9600

• 19200 (default)

• 57600

• 115200

For information about the options that you can use with the command show, see help show.

SEE ALSO

modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-12-02 sys console(1)

sys core

NAME

enhancedcorefiles - Configures management of core files for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the core component within the sys module using the following syntax.

MODIFY

modify core

options:

tmm-manage [true|false]

tmm-max [1-100]

tmm-action [skip|rotate]

mcpd-manage [true|false]

mcpd-max [1-100]

mcpd-action [skip|rotate]

bigd-manage [true|false]

bigd-max [1-100]

bigd-action [skip|rotate]

retention [1-365]

edit core

options:

all-properties

non-default-properties

DISPLAY

list core

list core [option]

show running-config sys core
show running-config sys core [option]
options:
all-properties
non-default-properties
one-line

DESCRIPTION

You can use this component to configure the creation of core files that the system will use to generate core files.

EXAMPLES

```
modify sys core tmm-manage true
```

Configures the TMOS system to manage tmm corefiles.

OPTIONS

process-manage
Whether or not this process' core files are managed by this component.

process-max
Maximum number of core files that may exist in /shared/cores for this process before the process-action is taken.

process-action
Action to take when number of core files is at its maximum. Possible actions are "skip" which will skip creation of the core file. "rotate" will delete the oldest core file before creation of the new core file.

retention
Integer representing number of days before deletion.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013, 2019. All rights reserved.

BIG-IP 2019-04-25 sys core(1)

sys cpu

NAME

cpu - Displays statistics about the Traffic Management Microkernel (TMM) service, specifically, CPU cycles.

MODULE

sys

SYNTAX

Display statistics for the cpu component within the sys module using the syntax in the following section.

DISPLAY

```
show cpu  
options:  
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)  
global
```

DESCRIPTION

You can use the cpu component to display the CPU cycles for the system. You can also specify the unit value in which the system displays statistics.

EXAMPLES

```
show cpu
```

Displays TMM processor statistics in the system default units.

```
show cpu raw
```

Displays raw TMM processor statistics.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or

mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2015. All rights reserved.

BIG-IP 2015-07-22 sys cpu(1)

sys crypto acceleration-strategy

NAME

acceleration-strategy - Control crypto operations offloading on the BIG-IP(r) system (hybrid mode).

MODULE

sys crypto

SYNTAX

Configure offload strategy

LIST

list acceleration-strategy

DESCRIPTION

Display the configured setting.

EXAMPLES

list sys crypto acceleration-strategy

MODIFY

modify acceleration-strategy [fixed-ratio]

DESCRIPTION

Set the fixed-ratio policy setting as a value in the interval 0-1000. The value 1000 means that cryptographic operations will be processed by available hardware accelerator(s). A value 700 means that 70% of cryptographic operations will be processed by available hardware accelerator(s), while remaining 30% of operations will be processed on CPU.

Currently only the following operations are subject to this policy: ECDSA signing, RSA encrypt or verify, and ECDH.

EXAMPLES

modify sys crypto acceleration-strategy fixed-ratio 700

In the above example, 70% (700 out of 1000) of relevant cryptographic requests are processed by hardware accelerator(s), if available, while 30% (300 out of 1000) of relevant cryptographic requests will be sent to the CPU.

SHOW

show acceleration-strategy

DESCRIPTION

Display acceleration-strategy statistics.

EXAMPLES

show sys crypto acceleration-strategy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-10-19 sys crypto acceleration-strategy(1)

sys crypto allow-key-export

NAME

allow-key-export - Specifies whether or not to allow private key export.

LIST/MODIFY

list allow-key-export

modify allow-key-export

options:
value [enabled | disabled]

EXAMPLES

```
list sys crypto allow-key-export value
```

```
modify sys crypto allow-key-export value enabled
```

```
modify sys crypto allow-key-export value disabled
```

```
BIG-IP      2017-07-20      sys crypto allow-key-export(1)
```

sys crypto ca-bundle-manager

NAME

ca-bundle-manager - Certificate Authority (CA) certificate bundle manager on the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

A ca-bundle-manager manages cryptographic ca-bundles using the syntax given in the following sections.

CREATE/MODIFY

```
create ca-bundle-manager [name]
```

```
modify ca-bundle-manager [name]
```

options:

```
description [string]
```

```
exclude_bundle
```

```
[add | delete | replace-all-with ] {
```

```
[cert file obj] ...
```

```
}
```

```
exclude_url
```

```
[add | delete | replace-all-with ] {
```

```
[url] ...
```

```
}
```

```
include_bundle
```

```
[add | delete | replace-all-with ] {
```

```
[cert file obj] ...
```

```
}
```

```
include_url
```

```
[add | delete | replace-all-with ] {
```

```
[url] ...
```

```
}
```

```
proxy-server [ [hostname] | [ipv4] | [ipv6] ]
```

```
proxy-port [ port number ]
```

```
trusted-ca-bundle [certificate file object]
```

```
update-interval [days]
```

```
time-out [seconds]
```

```
update-now [yes | no]
```

LIST

```
list ca-bundle-manager [name]
```

options:

```
-hidden
```

DELETE

```
delete ca-bundle-manager [name]
```

DESCRIPTION

You can use the ca-bundle-manager component to automatically update and install CA-bundles on the system from two sources - local certificate file objects and remote URL resources, using set include/exclude operations.

The set include/exclude operations are equivalent to mathematical set addition/subtraction operations. For example, the user may use include-bundle and include-url options to combine CA-certificates from various sources, and use exclude-bundle and exclude-url options to remove certain CA-certificates from the final CA-bundle file. The generated CA-bundle file will be installed as a certificate-file-object on the system, and used as trusted CA-bundle by other modules. Additionally, the user may set the update frequency of the CA-bundle, or use web proxy for downloading the remote URL resources. By default, a newly created CA-bundle manager does not create or update the managed CA-bundle object unless it has a positive update interval or being explicitly told to do so by the update-now option. Additionally, the calculated CA-bundle must contain at least two CA certificates to be installed on the system.

EXAMPLES

```
modify sys crypto ca-bundle-manager bmgr include-bundle add { ca-bundle.crt } include-url add {  
https://ca.f5net.com/ca-bundle.crt } trusted-ca-bundle trusted-ca-chain.crt update-interval 30
```

Creates a ca-bundle-manager bmgr from two sources, one is a locally installed certificate file object ca-bundle.crt, and the other is from remote URL resource https://ca.f5net.com/ca-bundle.crt using trusted CA bundle . bmgr is refreshed from the two sources every 30 days.

modify sys crypto ca-bundle-manager bmgr update-now yes

Extending from above example, this command triggers an immediate update of the generated ca-bundle from its sources.

list sys crypto ca-bundle-manager bmgr -hidden

Shows all the properties of the ca-bundle-manager bmgr, including the hidden fields.

delete sys crypto ca-bundle-manager bmgr

Deletes the ca-bundle-manager bmgr from the system. Note that the generated ca-bundle certificate file object is not removed, and can still be used.

OPTIONS

description

Specifies user defined description.

include-bundle

Specifies a list of certificate file objects to include for generating the new ca-bundle.

include-url

Specifies a list of remote ca-bundles at the URLs to include for generating the new ca-bundle.

exclude-bundle

Specifies a list of certificate file objects to exclude from the new ca-bundle.

exclude-url

Specifies a list of remote ca-bundles at the URLs to exclude from the new ca-bundle.

partition Displays the administrative partition within which this ca-bundle-manager resides.

proxy-server Specifies the host name or IP address of the proxy server for accessing remote URL resources.

Only HTTP proxy is supported. Optional http:// may be prepended.

proxy-port Specifies the port number of the proxy server for accessing remote URL resources. Default is 3128.

trusted-ca-bundle

Specifies the trusted CA certificate bundle when downloading ca-bundles from the other URLs.

update-interval

Specifies the update interval in days to refresh the remote ca-bundles at the URLs. Default value is 0, which means the generated ca-bundle is not dynamically updated.

time-out

Specifies the time-out period in seconds to download the remote ca-bundles at the URLs. The value ranges between 1 and 3600 (1 hour). The default value is 8 seconds.

update-now

Specifies whether the ca-bundle-manager should immediately refresh its generated ca-bundle from all its sources and recalculate its certificate contents. The default value is no.

updated-by

Specifies a read-only attribute from which this ca-bundle-manager was last updated.

managed-bundle

Specifies a read-only attribute, which indicates the ca-bundle certificate file object name, managed by this ca-bundle-manager.

SEE ALSO

create, list, modify, delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2017-09-05 sys crypto ca-bundle-manager(1)

sys crypto cert-order-manager

NAME

cert-order-manager - Certificate order manager on the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

A cert-order-manager Manages the collection of Certificate Authority (CA) requirements for making certificate

orders using the syntax given in the following sections.

CREATE/MODIFY

create cert-order-manager [name]

modify cert-order-manager [name]

options:

app-service [[string] | none]

additional-headers [[string] | none]

authority [comodo | digicert | godaddy | symantec]

auto-renew [yes | no]

base-url [URL | none]

ca-cert [certificate file object]

client-cert [certificate file object | none]

client-key [certificate key file object | none]

client-key-passphrase [[string] | none]

edit-order-info

internal-proxy [internal proxy object]

login-name [[string] | none]

login-password [[string] | none]

order-info [string]

validity-days [days | none]

LIST

list cert-order-manager [name]

DELETE

delete cert-order-manager [name]

DESCRIPTION

cert-order-manager A component holds the Certificate Authority's (CA) specific requirements for making certificate orders. The user needs to select a CA from the supported list, configure the necessary authentication information, and order the information specific to the selected CA.

EXAMPLES

```
create sys crypto cert-order-manager certmgr authority comodo login-name cert-admin@myorg.com login-password
default ca-cert ca-bundle.crt internal-proxy iproxy-caapi additional-headers "customerUri:myorg-auto-poc"
order-info "{ orgId 5678 serverType -1 certType 136 }"
```

Creates a certificate order manager certmgr for certificate authority comodo. For CA account login authentication username cert-admin@myorg.com and password default is used. ca-bundle.crt is used for authenticating a TLS connection to a CA server and validating the certificate issued by the CA. customerUri:myorg-auto-poc provides customer Uri issued by comodo for the certificate requesting organization. In order info { orgId 5678 serverType -1 certType 136 } organization identity orgId 5678 is provided by comodo, and certType 136 is the certificate product type offered by comodo for the organization.

```
list sys crypto cert-order-manager certmgr
```

Shows all the properties of the cert-order-manager certmgr.

```
delete sys crypto cert-order-manager certmgr
```

Deletes the cert-order-manager certmgr from the system.

OPTIONS

additional-headers

Specifies additional headers required for the certificate authority with expected format "key:value,...". For example: (comodo) "customerUri:mycomp-auto-poc"

authority

Specifies a certificate authority.

auto-renew

Enable/Disable the certificate automatic renewals. By default, the automatic certificate renewal is enabled.

base-url

Specifies the base-url for reaching the CA. This is an optional field which gets populated with default values for a specific certificate authority.

ca-cert

Specifies the CA certificate to be used for authenticating the TLS connection with the CA server. ca-cert is also used for validating an issued certificate from CA before accepting into the system.

client-cert

Specifies the client authentication certificate used for accessing the CA account. This is a required field for certain CA accounts.

client-key

Specifies the client authentication key used for accessing the CA account. This is a required field for certain CA accounts.

client-key-passphrase

Specifies the optional key passphrase required for decrypting the client-key.

edit-order-info

Provides an editor for creating and modifying the order-info configuration. This should be the last property since selecting save and exit from the editor automatically submits the configuration.

internal-proxy

Specifies the internal proxy object that should be used for reaching the CA server.

login-name

Specifies the login name for accessing the CA account. This is a required field for certain CA accounts.

login-password

Specifies the login password for accessing the CA account. This is a required field for certain CA accounts.

order-info

Specifies a string containing necessary information for making certificate orders with CA. Format and fields of order-info varies with the CA.

validity-days

Specifies certificate validity in days. The default value is 365 days.

SEE ALSO

create, list, modify, delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-12-06 sys crypto cert-order-manager(1)

sys crypto cert-validation-response ocs

NAME

ocsp - Manages the OCSF response of a certificate when it has OCSF validation enabled.

MODULE

sys crypto cert-validation-response

SYNTAX

Use the ocs component within the sys.crypto.cert-validation-response module to manage the cached OCSF response of a given certificate.

DELETE

delete ocs

options:

certificate [name]

DESCRIPTION

You can use the ocs component to delete the cached OCSF response of a given certificate. The response is automatically re-fetched, and the certificate status is updated subsequently.

OPTIONS

certificate

Specifies the name of the certificate whose OCSF response you want to delete.

SEE ALSO

delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2016-09-14 sys crypto cert-validation-response ocs(1)

sys crypto cert-validator cri

NAME

cri - Configures a CRL cert-validator.

MODULE

sys crypto cert-validator

SYNTAX

Configure the `crl` component using the syntax shown in the following sections.

CREATE/MODIFY

```
create crl [name]
modify crl [name]
options:
  internal-proxy [none | [string] ]
  strict-revocation-check [disabled | enabled]
reset-stats crl [name]
```

DISPLAY

```
list crl [name]
show crl [name]
```

DELETE

```
delete server-ssl [all | [name]]
options:
  recursive
```

DESCRIPTION

You can use the `crl` component to create, modify, display or delete a custom CRL cert-validator.

EXAMPLES

```
create cert-validator crl my_crl internal-proxy my_intp
```

Creates a CRL cert-validator named `my_crl` using the internal proxy named `my_intp`.

OPTIONS

internal-proxy

Specifies the internal proxy to define the route for reaching CRL distribution points to fetch CRL files.

strict-revocation-check

If enabled, the certificate status won't be checked until the CRL fetching and caching is complete. If disabled, the certificate status will be immediately considered as unknown if the CRL fetching and caching is not complete yet. The default value is disabled.

SEE ALSO

sys crypto cert-validator ocp, sys internal-proxy

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2017. All rights reserved.

BIG-IP 2017-12-11 sys crypto cert-validator crl(1)

sys crypto cert-validator ocp

NAME

ocp - Configuration of the OCP cert-validator.

MODULE

sys crypto

SYNTAX

Configure the `ocp` component within the `sys.crypto.cert-validator.ocp` module using the syntax shown in the following sections. This object is associated with a certificate object to enable an OCP request for updating the certificate status.

CREATE/MODIFY

```
create ocp [name]
modify ocp [name]
options:
  cache-error-timeout [integer]
  cache-timeout [indefinite | [integer] ]
  concurrent-connections-limit [integer]
  clock-skew [integer]
  description [string]
  dns-resolver [name]
  proxy-server-pool [name]
  responder-url [none | [string] ]
  route-domain [name]
  sign-hash [sha1 | sha256]
```

signer-cert [name]
signer-key [name]
signer-key-passphrase [none | [string]]
status-age [integer]
strict-resp-cert-check [disabled | enabled]
timeout [indefinite | [integer]]
trusted-responders [none | [name]]

DISPLAY

list ocsf [name]

DELETE

delete [all | [name]]

options:

recursive

DESCRIPTION

You can use the ocsf component to create, modify, display or delete a custom OCSF cert-validator.

The OCSF cert-validator is associated with a certificate object.

EXAMPLES

create cert-validator my_ocsf dns-resolver name

Creates an OCSF cert-validator named my_ocsf using the DNS resolver specified by name.

OPTIONS

cache-error-timeout

Specifies the lifetime of an error response in the cache, in seconds. The default value is 3600 seconds.

cache-timeout

Specifies the lifetime of the OCSF response in the cache, in seconds. The actual time period for which the response is cached is the minimum of the response validity period and the cache-timeout. The default value is indefinite, indicating that the response validity period takes precedence.

concurrent-connections-limit

Specifies the maximum number of connections per second allowed for the OCSF cert-validator.

clock-skew

Specifies the tolerable absolute difference in the clocks of the responder and the BIG-IP, in seconds. The default value is 300.

description

User defined description.

dns-resolver

Specifies the DNS resolver object used for fetching the OCSF response.

partition

Displays the administrative partition within which this validator resides.

proxy-server-pool

Specifies the proxy server pool used for fetching the OCSF response.

responder-url

Specifies the absolute URL that overrides the OCSF responder URL obtained from the certificate's AIA extension(s). This should be an HTTP-based URL.

route-domain

Specifies the route domain for fetching an OCSF response using HTTP forward proxy.

sign-hash

Specifies the hash algorithm used for signing the OCSF request. The default value is sha256.

signer-cert

Specifies the certificate corresponding to the key used for signing the OCSF request.

signer-key

Specifies the key used for signing the OCSF request.

signer-key-passphrase

Specifies the passphrase of the key used for signing the OCSF request.

status-age

Specifies the maximum allowed lag time for the 'thisUpdate' time in the OCSF response that the BIG-IP accepts. If this maximum is exceeded, the response is dropped. If this value is set to 0, this validation is skipped. The default value is 86400 seconds.

strict-resp-cert-check

If enabled, the responder's certificate is checked for an OCSF signing extension. The default value is disabled.

timeout

Specifies the time interval (in seconds) that the BIG-IP waits for before ending the connection to the OCSF responder. The default value is 8.

trusted-responders

Specifies the certificates used for validating the OCSF response when the responder's certificate has been omitted from the response.

SEE ALSO

create, delete, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2017-01-20 sys crypto cert-validator ocsf(1)

sys crypto cert

NAME

cert - Manage cryptographic certificates on the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

Manage cryptographic certs using the syntax in the following section.

CREATE

create cert [name]

options:

city [string]
common-name [string]
consumer
[enterprise-manager | iquery | iquery-big3d | ltm | webserver]
country [string]
email-address [string]
key [string]
lifetime [days]
organization [string]
ou [string]
state [string]
subject-alternative-name [string]

INSTALL

install cert [name]

options:

cert-validation-options [none | ocsf]
cert-validators [none | [cert_validator_name]]
consumer
[enterprise-manager | iquery | iquery-big3d | ltm | webserver]
from-editor
from-local-file [filename]
from-url [URL]
issuer-cert [none | [issuer_cert_name]]
no-overwrite

MODIFY

modify cert [name]

options:

cert-validation-options [none | ocsf]
cert-validators [none | [cert_validator_name]]
issuer-cert [none | [issuer_cert_name]]

DELETE

delete cert [name]

DESCRIPTION

You can use the cert component to create, install, and delete cryptographic certificates, and bundles.

EXAMPLES

```
create cert example key testkey.key common-name "My Company Inc." country "US"
```

Generates a self signed certificate named "example.crt". A key with the specified name "testkey.key" in this case must be installed on the system in order for this operation to succeed. The cert extension (".crt") will be appended to the created cert name if it is not already provided in the name.

```
create cert /myfolder/example key testkey.key common-name "My Company Inc." country "US"
```

Similar to above, but creates the cert "example.crt" in the folder "/myfolder" instead of the default "/Common". The specified folder "/myfolder" must already exist in order for this operation to succeed.

create cert server2 key server2.key common-name "My Company Inc." country "US" consumer webserver

Generates a self-signed certificate named server2.crt. The consumer attribute, "webserver", is used to cause the files to be placed directly in the path which can be found by the BIG-IP system httpd. A pre-existing key named "server2.key" must exist in the web server's key path in order for this operation to succeed. Please note that for non LTM consumer's key and cert names must be the same.

install cert example from-editor

Opens an interactive editor session into which can be pasted a certificate for import into the BIG-IP system. A certificate file-object will be created with the name example which contains the contents saved from the editor session.

install cert example from-local-file /tmp/example.crt

Obtains a certificate from the file located at /tmp/example.crt.

install cert example from-url http://example.com/example.crt

Obtains a certificate from a remote host, based on the URI specified.

modify sys crypto cert leaf.crt issuer-cert issuer.crt cert-validators add { my_ocsp1 } cert-validation-options { ocsp }

Assigns issuer certificate issuer.crt to the certificate leaf.crt, associates the OCSP certificate validator my_ocsp to the certificate, and enables the OCSP certificate validator for this certificate.

delete cert example.crt

Deletes the certificate "example.crt" from the system.

OPTIONS

cert-validation-options

Specifies the option used for validating the certificate status.

cert-validators

Specifies the name of the cert-validators used for validating the certificate status. Each cert-validation type can only have one cert-validator.

city Specifies the x509 city field to be used in creation of the certificate.

common-name

Specifies the x509 common-name to be used in creation of the certificate.

consumer

Specifies the system component by which a certificate will be consumed. The default behavior is to create file-objects for use by ltm components. This is the same as specifying "ltm" for this property. If a component other than "ltm" is specified then files will be installed/created in locations where the specified components can find them. For example, for component "webserver", certificates will be placed in the webserver's ssl directories.

country

Specifies the x509 country to be used in creation of the certificate. The country must be a 2 letter country code.

email-address

Specifies the x509 email-address to be used in creation of the certificate.

fingerprint

Displays the SHA-256 fingerprint of the certificate.

from-editor

Specifies that the certificate should be obtained from a text editor session. This allows certificates to be imported via cut-n-paste from another location as long as they are in a text representation.

from-local-file

Specifies a local file path from which a certificate is to be copied.

from-url

Specifies a URI which is to be used to obtain a certificate for import into the system.

The URL syntax is protocol dependent. Supported schemes are "HTTP", "HTTPS", "FTP", "FTPS" & "FILE."

issuer-cert

Specifies the name of the issuer certificate for this certificate.

no-overwrite

Specifies option of not overwriting a certificate if it is in the scope.

key Specifies a key from which a certificate should be generated when using the create command.

organization

Specifies the x509 organization to be used in creation of the certificate.

ou Specifies the x509 organizational unit to be used in creation of the certificate.

state
Specifies the x509 state or province of the certificate.

subject-alternative-name
Specifies standard X.509 extensions as shown in RFC 2459. Allowed values e.g. DNS:example.com, IP:192.168.1.1, IP:12:34, email:user@example.com, URI:http://www.example.com

SEE ALSO
create, install, modify, delete, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2017-05-01 sys crypto cert(1)

sys crypto check-cert

NAME
check-cert - Examines certificates and displays or logs any that have expired on the BIG-IP(r) system.

MODULE
sys crypto

SYNTAX
Run a check on the expiration date of LTM certificates, in the sys crypto module by using the syntax below.

RUN
run check-cert [certificate-file-name]
options:
ignore-large-cert-bundles [enabled | disabled]
log [enabled | disabled]
stdout [enabled | disabled]
verbose [enabled | disabled]

DESCRIPTION
You can use the check-cert command to check the expiration date of certificate(s) and print the results to the screen and/or log them to /var/log/ltm.

OPTIONS
ignore-large-cert-bundles
Specifies whether or not to ignore large certificate bundles which contain more than 20 certificates. By default it will not be ignored, i.e., it will still check every certificate bundle if this option is not specified.

log Specifies whether results should be logged or not. By default they will be logged.

stdout
Specifies whether results should be printed to STDOUT or not. By default they will be printed.

verbose
Specifies whether verbose output should be emitted or not, such as information about all certificates being checked rather than just those which return unfavorable results. By default verbose output is disabled.

EXAMPLES
run check-cert

Checks all certificate file-objects known by MCPD, and displays information about any certificates which have expired or which are close to expiration. By default this information is printed to the screen and logged to /var/log/ltm.

run check-cert default.crt

Runs the check on the specific certificate "default.crt"

run check-cert verbose

Displays expiration information about all certificates, not just those that have expired or have impending expirations.

run check-cert ignore-large-cert-bundles enabled

Ignore the certificate bundles with large size (the ones containing more than 20 certificates).

run check-cert log disabled

Prints the results to screen but does not log them.

run check-cert stdout disabled

Logs the results to /var/log/ltm, but does not print them to the screen.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys crypto check-cert(1)

sys crypto client

NAME

client - Manage remote crypto clients.

MODULE

sys crypto

SYNTAX

Manage crypto clients using the syntax in the following section.

CREATE/MODIFY

create client [name]

modify client [name]

options:

addr [ip address]

connection-reset

[disabled | enabled]

heartbeat [integer]

max-retries [integer | infinite]

port [integer]

profiles [add | delete | replace-all-with] { [profile_name ...] }

profiles [none]

req-timeout [integer]

retry-interval [integer]

DISPLAY

list client

list client [[name] | [globl] | [regex]] ...]

show client

show client [[name] | [globl] | [regex]] ...]

DELETE

delete client [name]

DESCRIPTION

You can use the client component to manage remote crypto clients.

EXAMPLES

create client example addr 10.1.1.1 port 12100 profiles add { serverssl tcp }

Creates a remote crypto client named "example" that will use a remote crypto server with the IP address "10.1.1.1" on port "12100". The remote crypto client will use SSL over a TCP connection to communicate with the remote crypto server.

OPTIONS

addr Specifies the IP address of the remote crypto server.

connection-reset

Resets the connection to the remote crypto server.

(enabled | disabled)

Specifies the state of the remote crypto client. The default value is enabled.

heartbeat

Specifies the number of seconds to wait before sending a heartbeat request. A value of 0 disables the sending of heartbeat requests. The default value is 30 seconds.

max-retries

Specifies the maximum number of times to retry connecting to the remote crypto server.

If the maximum retries value is infinite, the crypto client retries connecting until a connection is made.

The default value is infinite.

port Specifies the port used by the remote crypto server.

profiles
Specifies a list of profiles that the remote crypto client will use to communicate with remote the remote crypto server.

req-timeout
Specifies the timeout in milliseconds for crypto requests to complete. The default value is 5000 milliseconds.

retry-interval
Specifies the interval in seconds between attempts to connect to the remote crypto server. The default value is 10 seconds.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2015-01-16 sys crypto client(1)

sys crypto crl

NAME

crl - Manage certificate revocation lists on the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

Manage crls using the syntax in the following section.

INSTALL

install crl [name]

options:

ca-file [filename]

consumer

[enterprise-manager | iquery | iquery-big3d | ltm | webserver]

from-editor

from-local-file [filename]

from-url [URL]

DELETE

delete crl [name]

DESCRIPTION

You can use the crl component to install, and delete certificate revocation lists. The file-objects created by these operations can be used in other BIG-IP system configuration blocks such as ssl profiles.

EXAMPLES

install crl example from-editor

Opens an interactive editor session into which can be pasted a crl for import into the BIG-IP system. A crl file-object will be created with the name example which contains the contents saved from the editor session.

install crl example from-local-file /tmp/example.crl

Obtains a crl from the file located at /tmp/example.crl and installs it as example.crl. The crl extension (".crl") will be appended to the installed crl name if it is not already provided in the name.

install crl /myfolder/myexample from-local-file /tmp/example.crl

Similar to above, but installs the crl "myexample.crl" in folder "/myfolder" instead of the default "/Common". The specified folder "/myfolder" must already exist in order for this operation to succeed.

install crl example from-url http://example.com/example.crl

Obtains a crl from a remote host, based on the URI specified.

delete crl example.crl

Deletes the certificate revocation list "example.crl" from the system.

OPTIONS

consumer

Specifies the system component by which the certificate revocation list will be consumed. The default behavior is to create file-objects for use by Itm components. This is the same as specifying "Itm" for this property. If a component other than "Itm" is specified then files will be installed/created into locations where the specified components can find them. For example, for component "webserver", crls will be placed in the webservers ssl directories.

from-editor

Specifies that the crl should be obtained from a text editor session. This allows crls to be imported via cut-n-paste from another location as long as they are in a text representation.

from-local-file

Specifies a local file path from which the crl is to be copied.

from-url

Specifies a URI which is to be used to obtain the crl for import into the configuration of the system.

The URL syntax is protocol dependent. Supported schemes are "HTTP", "HTTPS", "FTP", "FTPS" & "FILE."

SEE ALSO

create, install, delete, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-03-21 sys crypto crl(1)

sys crypto csr

NAME

csr - Manage cryptographic certificate signing requests on the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

Manage cryptographic CSRs using the syntax in the following section.

CREATE

create csr [name]

options:

challenge-password [string]

admin-email-address [string]

city [string]

common-name [string]

consumer

[enterprise-manager | iquery | iquery-big3d | Itm | webserver]

country [string]

email-address [string]

key [string]

organization [string]

ou [string]

state [string]

subject-alternative-name [string]

SHOW

show csr

LIST

list csr [name]

DELETE

delete csr [name]

DESCRIPTION

You can use the csr component to create, show, list and delete cryptographic certificate signing requests.

EXAMPLES

create csr example key testkey.key common-name "My Company Inc." country "US" challenge-password "abcd"

Generates a certificate signing request named "example.csr" with provided common-name, country and challenge-password attributes. A key with the specified name "testkey.key" in this case must be installed on the system in order for this operation to succeed. The csr extension (".csr") will be appended to the created csr name if it is not already provided in the name.

```
create csr /myfolder/example key testkey.key common-name "My Company Inc." country "US" challenge-password "abcd"
```

Similar to above, but creates the csr "example.csr" in the folder "/myfolder" instead of the default "/Common". The specified folder "/myfolder" must already exist in order for this operation to succeed.

```
create csr server2 key server2.key common-name "My Company Inc." country "US" consumer webserver
```

Generates a certificate signing request named "server2.csr". The consumer attribute, "webserver", is used to cause the files to be placed directly in the path which can be found by the BIG-IP system httpd. A pre-existing key named "server2.key" must exist in the web server's key path in order for this operation to succeed.

```
show csr
```

Shows the number of certificate signing requests installed in the system.

```
list csr example.csr
```

Lists all details of the certificate signing request "example.csr". A csr with the specified name "example.csr" in this case must already be installed on the system in order for this operation to succeed. Because only one certificate signing request name is specified in the list command, it will also display the contents of the certificate signing request file.

```
list csr example1.csr example2.csr
```

Lists all details of the certificate signing requests "example1.csr" and "example2.csr". Because more than one certificate signing request name is specified in the list command, it will not display the contents of the certificate signing request files.

```
list csr
```

Lists details of all certificate signing requests that are configured in the system. This command does not display the contents of the certificate signing request files.

```
delete csr example.csr
```

Deletes the certificate signing request "example.csr" from the system.

OPTIONS

challenge-password

Specifies the PKCS#9 challenge-password field to be used in creation of the certificate signing request.

admin-email-address

Specifies the administrator email-address to be used in creation of the certificate signing request.

city Specifies the x509 city field to be used in creation of the certificate signing request.

common-name

Specifies the x509 common-name to be used in creation of the certificate signing request.

consumer

Specifies the system component by which a certificate signing request will be consumed. The default behavior is to create file-objects for use by ltm components. This is the same as specifying "ltm" for this property. If a component other than "ltm" is specified then files will be installed/created in locations where the specified components can find them. For example, for component "webserver", certificate signing requests will be placed in the webserver's ssl directories.

country

Specifies the x509 country to be used in creation of the certificate signing request. The country must be a 2 letter country code.

email-address

Specifies the x509 email-address to be used in creation of the certificate signing request.

key Specifies a key from which a certificate signing request should be generated when using the create command.

organization

Specifies the x509 organization to be used in creation of the certificate signing request.

ou Specifies the x509 organizational unit to be used in creation of the certificate signing request.

state

Specifies the x509 state or province to be used in creation of the certificate signing request.

subject-alternative-name

Specifies standard X.509 subject alternative extensions as shown in RFC 2459 to be used in creation of the certificate signing request. Examples of allowed types are : DNS:example.com, IP:192.168.1.1, IP:12:34, email:user@example.com, URI:http://www.example.com

SEE ALSO

create, show, list, delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2015. All rights reserved.

BIG-IP 2016-07-20 sys crypto csr(1)

sys crypto encrypted-attributes

NAME

encrypted-attributes - Displays all objects in the MCP binary database that contain items encrypted with the master key for the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

Display a list of objects with attributes encrypted with the master-key in the binary database using the syntax in the following section.

DISPLAY

show encrypted-attributes

options:

field-fmt

DESCRIPTION

You can use the encrypted-attributes command to help diagnose master-key problems in the BIG-IP(r) system. Only users with the Administrator role or the Resource Administrator role can view the list of objects with encrypted attributes.

EXAMPLES

show encrypted-attributes

Displays, in a table, information about the system's objects with encrypted-attributes including type, object name, attribute name, and whether or not the encrypted value in the binary database can be decrypted with the current master-key.

show master-key field-fmt

Displays, in field format, information about the system's objects with encrypted-attributes.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2016-03-14 sys crypto encrypted-attributes(1)

sys crypto fips by-handle

NAME

by-handle - NOTICE: The use of FIPS 140 key handles has been deprecated in TMOS 12.1.0 in favor of key IDs

MODULE

sys crypto fips

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012, 2015. All rights reserved.

BIG-IP 2015-09-24 sys crypto fips by-handle(1)

sys crypto fips external-hsm

NAME

external-hsm - Configures parameters for external HSM FIPS hardware.

MODULE

sys crypto fips

DESCRIPTION

You can use the external-hsm command to set parameters about the HSM vendor name and the password to login to the external HSM hardware.

SYNTAX

Configures FIPS external-hsm within the sys crypto fips module using the syntax in the following section.

CREATE

```
create external-hsm
modify external-hsm vendor [thales | safenet | auto | none]
modify external-hsm password [password]
modify external-hsm pkcs11-lib-path [path to pkcs#11 library provided by the vendor]
modify external-hsm num-threads [no. of threads]
```

DISPLAY

```
list external-hsm
list external-hsm vendor
list external-hsm password
list external-hsm pkcs11-lib-path
list external-hsm num-threads
```

DELETE

```
delete external-hsm
```

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

BIG-IP 2018-02-08 sys crypto fips external-hsm(1)

sys crypto fips key

NAME

key - Manage FIPS keys

MODULE

sys crypto fips

SYNTAX

Manage cryptographic keys within the sys crypto fips module using the syntax in the following section.

SHOW

```
show key [key ID]
options:
  field-fmt
  all-properties
  include-public-keys
```

DELETE

```
delete key [key ID]
```

DESCRIPTION

You can use the key component to show and delete cryptographic keys contained in the FIPS hardware.

EXAMPLES

```
show key
```

Displays the list of all private keys stored in the FIPS hardware and their meta-data.

show key bef8221fd25a27cda51b3904bc2f5fb8

Displays information specifically about the FIPS key with the key ID "bef8221fd25a27cda51b3904bc2f5fb8".

show key field-fmt

Displays, in field format, information about private keys stored in the FIPS hardware.

show key all-properties

Displays all information about the FIPS contained private keys, including: handle, a numerical value used by the FIPS hardware to identify individual keys; modulus-length, the cryptographic modulus length of the key; and modulus, the modulus associated with the key, displayed as a string of hex octets separated by colons.

show key include-public-keys

Displays the list of all private and public keys stored in the FIPS hardware and their meta-data. Note that public keys are not displayed by default and need not exist for normal operation of FIPS hardware.

delete key bef8221fd25a27cda51b3904bc2f5fb8

Deletes the FIPS key with the key ID "bef8221fd25a27cda51b3904bc2f5fb8" from the system.

OPTIONS

include-public-keys

Specifies that public keys should be selected for output in addition to private.

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2015-09-18 sys crypto fips key(1)

sys crypto key

NAME

key - Manage cryptographic keys and related objects on the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

Manage cryptographic keys and related objects of the sys crypto module using the syntax in the following section.

CREATE

create key [name]

options:

challenge-password [string]

admin-email-address [string]

city [string]

common-name [string]

consumer

[enterprise-manager | iquery | iquery-big3d | ltm | webserver]

country [string]

curve-name [prime256v1 | secp384r1 | secp521r1]

email-address [string]

key-size [512 | 1024 | 2048 | 4096]

key-type [dsa-private | ec-private | rsa-private]

lifetime [days]

organization [string]

ou [string]

passphrase [passphrase]

prompt-for-password

security-type [fips | normal | password | nethsm]

state [string]

subject-alternative-name [string]

cert-order-manager [add | delete | modify | replace-all-with] {

```

    [ [name] ] {
options:
check-status [yes | no]
order-id [string | none]
order-passphrase [string | none]
order-type [cancel | new | renew | revoke]
revoke-reason [AACompromise | affiliationChanged | cessationOfOperation | removeFromCRL | unspecified | CACompromise
}
}

```

MODIFY

modify key [name]

options

```

cert-order-manager [add | delete | modify | replace-all-with] {
    [ [name] ] {

```

options:

```

check-status [yes | no]
order-id [string | none]
order-passphrase [string | none]
order-type [cancel | new | renew | revoke]
revoke-reason [AACompromise | affiliationChanged | cessationOfOperation | removeFromCRL | unspecified | CACompromise
}
}

```

SHOW

show key

show key [name | none] cert-order-manager

LIST

list key

list key [name]

INSTALL

install key [name]

options:

```

consumer
    [enterprise-manager | iquery | iquery-big3d | itm | webserver]
from-editor
from-local-file [filename]
from-url [URL]
from-nethsm
no-overwrite

```

DELETE

delete key [name]

DESCRIPTION

You can use the key component to create, show, list, install, and delete cryptographic keys and associated cryptographic objects. The file-objects created by these operations can be used in other BigIP configuration blocks such as ssl profiles.

EXAMPLES

create key mykey

Generates a 2048-bit (default-sized) RSA key file object named "mykey.key". The appropriate extension will be added to the generated key/cert if not already a part of the provided name.

create key mykey key-type ec-private curve-name prime256v1

Generates a prime256v1 curve name EC private key file object named "mykey.key". The appropriate extension will be added to the generated key/cert if it is not already a part of the provided name.

create key /myfolder/mykey

Similar to the above action except it creates the key "mykey.key" in the folder "/myfolder" instead of the default "/Common". The specified folder "/myfolder" must already exist in order for this operation to succeed.

create key example gen-cert gen-csr common-name "My Company Inc." country "US"

Generates a 2048-bit (default-sized) RSA key file object named "example.key" and a self signed certificate named "example.crt". Also, a certificate signing request will be printed to the console for use in obtaining a signed certificate from a certificate authority, if desired.

create key my gen-cert gen-csr prompt-for-password common-name "My Company Inc." country "US"

Similar to the above action when creating key "my.key" except it also prompts for a password to be used as a challenge password in the certificate authority signing procedure.

create key server2 gen-cert gen-csr common-name "My Company Inc." country "US" consumer webserver

Generates a key and self signed certificate identified by server2. The consumer attribute, "webserver", is used to cause these files to be placed directly in the paths which can be found by the BigIP's httpd.

create key server gen-csr common-name "My Company Inc." country "US" cert-order-manager add { certmgr { order-type new } }

Generates a key and CSR identified by server. Associates cert-order-manager object "certmgr" with the key and

makes a "new" certificate order to the CA.

show key

Shows the number of keys installed in the system.

show key cert-order-manager

Shows certificate order statistics if a cert-order-manager object is associated with key.

list key example.key

Lists all details of the key named "example.key". A key with the specified name "example.key" in this case must already be installed on the system in order for this operation to succeed.

list key

Lists all details of all keys installed in the system.

install key example from-editor

Opens an interactive editor session into which it a key for import into the BigIP system can be pasted. A key file-object will be created with the name example which contains the contents saved from the editor session.

install key example from-local-file /tmp/example.key

Obtains a key from the file located at /tmp/example.key.

install key example from-url http://example.com/my.key

Obtains a key from a remote host that is based on the URI specified.

delete key example.key

Deletes the key "example.key" from the system.

OPTIONS

challenge-password

Specifies the challenge password to create the certificate request key.

admin-email-address

Specifies the administrator email-address to be used in creation of the certificate request associated with the given key.

city Specifies the x509 city field to be used in creation of the certificate associated with the given key.

common-name

Specifies the x509 common-name to be used in creation of the certificate associated with the given key.

consumer

Specifies the system component by which a key and/or associated cryptographic file will be consumed. The default behavior is to create file-objects for use by ltm components. This is the same as specifying "ltm" for this property. If a component other than "ltm" is specified then files will be installed/created into locations where the specified components can find them. For example, for component "webserver", keys and certs will be placed in the webserver's ssl directories.

country

Specifies the x509 country to be used in creation of the certificate associated with the given key. The country must be a 2 letter country code.

curve-name

Specifies the curve name to be used in creation of the elliptic curve (EC) key. This option only applies when generating EC keys. The default value is prime256v1.

email-address

Specifies the x509 email-address to be used in creation of the certificate associated with the given key.

from-editor

Specifies that the key should be obtained from a text editor session. This allows keys to be imported via cut-n-paste from another location as long as they are in a text representation.

from-local-file

Specifies a local file path from which a key is to be copied.

from-url

Specifies a URI which is to be used to obtain a key for import into the configuration of the system.

The URL syntax is protocol dependent. Supported schemes are "HTTP", "HTTPS", "FTP", "FTPS" & "FILE."

from-nethsm

Specifies an option to import an existing key from network-HSM to BIG-IP config. The key label is specified as the key name to identify the key to be imported from network-HSM.

no-overwrite

Specifies the option for not overwriting a key if it is in the scope.

gen-certificate

Specifies that in addition to generating a key, a self-signed certificate will also be created. If this option is specified then x509 attributes should also be specified. Minimally, you must also specify a common-name.

gen-csr

Specifies that a certificate signing request should be generated along with the key. The CSR will be displayed to the terminal for the purposes of use in getting a certificate signed by an outside authority. X509 attributes must also be specified.

key-size

Specifies the size, in bits, of the key to be generated. This option does not apply when generating EC keys.

key-type

Specifies the type of cryptographic key to be generated. Default is rsa-private.

lifetime

Specifies the certificate life time to be used in creation of the certificate associated with the given key.

organization

Specifies the x509 organization to be used in creation of the certificate associated with the given key.

ou Specifies the x509 organizational unit to be used in creation of the certificate associated with the given key.

prompt-for-password

Specifies that a password should be prompted for and then used as a challenge password in generation of the CSR (Certificate Signing Request).

security-type

Specifies the level of security used in storing the key in question. For example, a security-type of FIPS means that the key should be stored on a FIPS card if one is available.

state

Specifies the x509 state or province of the certificate associated with the given key.

passphrase

Specifies an optional passphrase with which the key has been protected. It may be used by consumers of the key in the data-plane or control-plane to decrypt it.

subject-alternative-name

Specifies standard X.509 extensions as shown in RFC 2459. Allowed values e.g. DNS:example.com, IP:192.168.1.1, IP:12:34, email:user@example.com, URI:http://www.example.com

cert-order-manager

Specifies an optional cert-order-manager to be associated with the key.

check-status

Specifies that it checks the status of a certificate order. This command triggers an immediate status check query with CA for a current pending certificate order.

order-id

Specifies the order id for a certificate order. This order id is provided by the CA and the bigip stores it in the order-id field. Order id is required for certificate renewal and revoke. If the first certificate was not originally ordered from the bigip, the user needs to enter the order-id manually before making a certificate renewal or revoke.

order-passphrase

Specifies the order challenge passphrase. This is a CA specific requirement. Some CA's require a challenge passphrase for making a certificate order.

order-type

Specifies the type of certificate order to authority.

new : Make a new certificate order to the CA.

renew : Make a certificate renewal order to the CA.

revoke : Make a certificate revoke order to the CA.

cancel : Tries to cancel the previous certificate order.

revoke-reason

Specifies the reason for certificate revoke.

SEE ALSO

create, install, show, list, delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015-2018. All rights reserved.

sys crypto master-key

NAME

master-key - Displays the configuration of the master key for the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

Display the configuration of the master-key component within the sys crypto module using the syntax in the following section.

DISPLAY

show master-key
options:
field-fmt

MODIFY

modify master-key
options:
prompt-for-password

RUN

run master-key diagnostic

DESCRIPTION

You can use the master-key command to manipulate the system master key. Users with the Administrator role or the Certificate Manager role can set the key to a value of their choosing by using the 'prompt-for-password' option during a modify operation. All other roles, including Resource Administrators, are prohibited from setting the master key.

Use the 'diagnostic' option of the run command to test the key integrity.

EXAMPLES

show master-key

Displays, in a table, information about the system's master key.

show master-key field-fmt

Displays, in field format, information about the system's master key.

run master-key diagnostic

Loads the device key. Uses the device key to decrypt the master key file to test the integrity of the keys. On success, there is no output. There will be a response only if there is an error.

modify master-key prompt-for-password

Create a master-key based on a word or phrase of your choosing. You can use this to manually synchronize several devices without having to copy keys between them.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys crypto master-key(1)

sys crypto pkcs12

NAME

pkcs12 - Install pkcs12 keys and certificates on the BIG-IP(r) system.

MODULE

sys crypto

SYNTAX

Install keys and certificates from pkcs12 files using the syntax in the following section.

INSTALL

install pkcs12 [name]

options:

consumer

[enterprise-manager | iquery | iquery-big3d | ltm | webserver]

from-local-file [filename]

from-url [URL]

key-passphrase

key-security-type

[fips | password | normal]

passphrase [passphrase]

no-overwrite

DESCRIPTION

You can use the pkcs12 component to install cryptographic keys and certificates from pkcs12 formatted files. The file-objects created by these operations can be used in other BigIP configuration blocks such as ssl profiles.

EXAMPLES

install pkcs12 example from-local-file /tmp/example.p12

Obtains a pkcs12 from the file located at /tmp/example.p12, and installs the key and certificate from that file as file-objects named "example.key" and "example.crt" respectively.

install pkcs12 /myfolder/example from-local-file /tmp/example.p12

Similar to above, but installs the key "example.key" and cert "example.crt" in folder "/myfolder" instead of the default "/Common". The specified folder "/myfolder" must already exist in order for this operation to succeed.

install pkcs12 example prompt-for-password from-local-file /tmp/example.p12

Same as above but also prompts for a password which is to be used to decrypt the pkcs12 file.

install pkcs12 my from-url http://example.com/my.p12

Obtains a pkcs12 file from a remote host, based on the URL specified.

install pkcs12 server consumer webserver from-local-file /tmp/example.p12

Obtains a pkcs12 file from /tmp/example.p12 and installs the key and certificate from that file as file-objects that can be used by the "webserver". The consumer attribute, "webserver", is used to cause these files to be placed directly in the paths which can be found by the BigIP's httpd.

OPTIONS

consumer

Specifies the system component by which a key and associated certificate from a PKCS12 file will be consumed. The default behavior is to create file-objects for use by ltm components. This is the same as specifying "ltm" for this property. If a component other than "ltm" is specified then files will be installed/created into locations where the specified components can find them. For example, for component "webserver", keys and certs will be placed in the webserver's ssl directories.

from-local-file

Specifies a local file path from which the contents of the PKCS12 are to be read.

from-url

Specifies a URI which is to be used to obtain a PKCS12 for import into the configuration of the system.

The URL syntax is protocol dependent. Supported schemes are "HTTP", "HTTPS", "FTP", "FTPS" & "FILE."

key-passphrase

Specifies the passphrase to be used to encrypt the key.

key-security-type

Specifies the security type of the key. Default is set to "normal".

passphrase

Specifies the passphrase to be used to decrypt the PKCS12 file.

no-overwrite

Specifies option of not overwriting key/certificate if they are in the scope.

SEE ALSO

install, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

sys crypto server

NAME

server - Manage remote crypto servers.

MODULE

sys crypto

SYNTAX

Manage crypto servers using the syntax in the following section.

CREATE/MODIFY

create server [name]

modify server [name]

options:

addr [ip address]

clients [add | delete | replace-all-with] { [ip_addr/prefixlen ...] }

clients [none]

[disabled | enabled]

port [integer]

profiles [add | delete | replace-all-with] { [profile_name ...] }

profiles [none]

DISPLAY

list server

list server [[[name] | [globl] | [regex]] ...]

show server

show server [[[name] | [globl] | [regex]] ...]

DELETE

delete server [name]

DESCRIPTION

You can use the server component to manage remote crypto servers.

EXAMPLES

```
create server example addr 10.1.1.1 port 12100 profiles add { clientssl tcp }
```

Creates a remote crypto server named "example" that will listen for remote crypto clients on IP address "10.1.1.1" and port "12100". The remote crypto server will use SSL over a TCP connection to communicate with remote crypto clients.

OPTIONS

addr Specifies the IP address of the remote crypto server.

clients

Specifies a list of allowed client IP addresses and subnets. An empty list allows all clients.

(enabled | disabled)

Specifies the state of the remote crypto server. The default value is enabled.

port Specifies the port used by the remote crypto server.

profiles

Specifies a list of profiles that the remote crypto server will use to communicate with remote crypto clients.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014, 2016. All rights reserved.

sys daemon-ha

NAME

daemon-ha - Configures high availability for a BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the daemon-ha component within the sys module using the syntax in the following sections.

MODIFY

modify daemon-ha [name]

options:

heartbeat [enabled | disabled]

heartbeat-action [go-offline | go-offline-downlinks-restart |

go-offline-restart | reboot | restart | restart-all]

running [enabled | disabled]

edit daemon-ha [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

DISPLAY

list daemon-ha

list daemon-ha [[name] | [glob] | [regex]] ...]

show running-config daemon-ha

show running-config daemon-ha [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

not-running-action

one-line

running-timeout

DESCRIPTION

You can use the daemon-ha component to configure the daemons on the system that handle high availability for the BIG-IP system.

EXAMPLES

modify daemon-ha bigd running disabled

Disables the bigd daemon.

list daemon-ha bigd running-timeout

Displays the running timeout of the bigd daemon.

OPTIONS

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

heartbeat

Specifies whether heartbeat monitoring is enabled for the specified daemon. If monitoring is enabled and the daemon does not maintain its heartbeat the action specified by the value of the heartbeat-action option is taken.

The default value is enabled for all daemons, except the named daemon, which is disabled by default.

heartbeat-action

Specifies the action the system takes if the specified daemon does not maintain its heartbeat.

The default value is dependent on the specified daemon, the most common default value is restart.

name Specifies a unique name for the component. This option is required for the command modify.

not-running-action

Specifies the action that the system takes if the daemon is not running. This option is read-only.

The default value is dependent on the specified daemon, the most common default value is go-offline-downlinks.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

running

Specifies whether the running-timeout and non-running-action options are enabled. The default value is dependent on the specified daemon, the most common default value is enabled.

Note: This feature is implemented only for the daemons: tmm, mcpd, bcm56xxd, bigd, gtmd, clusterd, tmrouted, bd, datasyncd and tmrouted.

running-timeout

Specifies the amount of time (in seconds) that must elapse before the specified daemon is considered to be not running. This option is read-only.

The default value is dependent on the specified daemon.

SEE ALSO

edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013. All rights reserved.

BIG-IP 2014-02-07 sys daemon-ha(1)

sys daemon-log-settings clusterd

NAME

clusterd - Changes the log-level of or displays information about the daemon clusterd.

MODULE

sys daemon-log-settings

SYNTAX

Configure the clusterd component within the sys daemon-log-settings module using the syntax in the following sections.

MODIFY

modify clusterd

options:

log-level [critical | debug | error | informational | notice | warning]

edit clusterd

options:

all-properties
non-default-properties

DISPLAY

list clusterd

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the clusterd component to change the level of the messages about the clusterd daemon that appear in the system logs. Additionally, you can display information about the daemon.

EXAMPLES

list clusterd

Displays information about the clusterd daemon.

modify clusterd log-level critical

Changes the level of the messages about the clusterd daemon that display in the system log to critical.

OPTIONS

log-level

Specifies the level of log messages for the specified daemon that you want to display in the system log.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-31 sys daemon-log-settings clusterd(1)

sys daemon-log-settings csyncd

NAME

csyncd - Changes the log-level of or displays information about the daemon csyncd.

MODULE

sys daemon-log-settings

SYNTAX

Configure the csyncd component within the sys daemon-log-settings module using the syntax in the following sections.

MODIFY

modify csyncd

options:

log-level [critical | debug | error | informational | notice | warning]

edit csyncd

options:

all-properties

non-default-properties

DISPLAY

list csyncd

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the csyncd component to change the level of the messages about the csyncd daemon that appear in the system logs. Additionally, you can display information about the daemon.

EXAMPLES

list csyncd

Displays information about the csyncd daemon.

modify csyncd log-level critical

Changes the level of the messages about the csyncd daemon that display in the system log to critical.

OPTIONS

log-level

Specifies the level of log messages for the specified daemon that you want to display in the system log.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-31 sys daemon-log-settings csyncd(1)

sys daemon-log-settings icr-eventd

NAME

icr-eventd - Changes the log level of the daemon icr-eventd.

MODULE

sys daemon-log-settings

SYNTAX

Configure the icr-eventd component within the sys daemon-log-settings module using the syntax in the following sections.

MODIFY

modify icr-eventd

options:

log-level [critical | debug | error | informational | notice | warning]

edit icr-eventd
options:
all-properties
non-default-properties

DISPLAY
list icr-eventd
options:
all-properties
non-default-properties
one-line

DESCRIPTION
You can use the icr-eventd component to change the ltm log level of the icr-eventd daemon.

EXAMPLES
list icr-eventd

Displays log level of the icr-eventd daemon.

modify icr-eventd log-level warning

Changes the log level of the icr-eventd daemon to warning.

OPTIONS
log-level
Specifies the level of log messages for the icr-eventd daemon in the ltm logs.

SEE ALSO
edit, list, modify, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-08-29 sys daemon-log-settings icr-eventd(1)

sys daemon-log-settings icrd

NAME
icrd - Changes or displays the audit level of the daemon icrd.

MODULE
sys daemon-log-settings

SYNTAX
Configure the icrd component within the sys daemon-log-settings module using the syntax in the following sections.

MODIFY
modify icrd
options:
audit [none | modifications | all]

edit icrd
options:
all-properties
non-default-properties

DISPLAY
list icrd
options:
all-properties
non-default-properties
one-line

DESCRIPTION
You can use the icrd component to change the audit level of the icrd daemon.

EXAMPLES
list icrd

Displays audit log level of the icrd daemon.

modify icrd audit all

Changes the audit level of the icrd daemon to all. This means that all commands executed by the icrd daemon get audited

OPTIONS

audit

Specifies the audit level of log messages for the icrd daemon in the audit logs.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2015-01-22 **sys daemon-log-settings icrd(1)**

sys daemon-log-settings lind

NAME

lind - Changes the log-level of or displays information about the daemon lind.

MODULE

sys daemon-log-settings

SYNTAX

Configure the lind component within the sys daemon-log-settings module using the syntax in the following sections.

MODIFY

modify lind

options:

log-level [critical | debug | error | informational | notice | warning]

edit lind

options:

all-properties

non-default-properties

DISPLAY

list lind

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the lind component to change the level of the messages about the lind daemon that appear in the system logs. Additionally, you can display information about the daemon.

EXAMPLES

list lind

Displays information about the lind daemon.

modify lind log-level critical

Changes the level of the messages about the lind daemon that display in the system log to critical.

OPTIONS

log-level

Specifies the level of log messages for the specified daemon that you want to display in the system log.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-31 **sys daemon-log-settings lind(1)**

sys daemon-log-settings mcpd

NAME

mcpd - Changes the log-level of or displays information about the daemon mcpd.

MODULE

sys daemon-log-settings

SYNTAX

Configure the mcpd component within the sys daemon-log-settings module using the syntax in the following sections.

MODIFY

modify mcpd

options:

audit [all | disabled | enabled | verbose]

log-level [alert | critical | debug | emergency | error | informational | notice | panic | warning]

edit mcpd

options:

all-properties

non-default-properties

DISPLAY

list mcpd

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the mcpd component to change the level of the messages about the mcpd daemon that appear in the system logs. Additionally, you can display information about the daemon.

EXAMPLES

list mcpd

Displays information about the mcpd daemon.

modify mcpd log-level critical

Changes the level of the messages about the mcpd daemon that display in the system log to critical.

OPTIONS

audit

Enables or disables auditing for the mcpd daemon, and specifies verbose or all as the auditing level. The default is disabled.

log-level

Specifies the level of log messages for the specified daemon that you want to display in the system log.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-31 sys daemon-log-settings mcpd(1)

sys daemon-log-settings tmm

NAME

tmm - Changes the log-level of or displays information about the Traffic Management Microkernel (tmm).

MODULE

sys daemon-log-settings

SYNTAX

Configure the tmm component within the sys daemon-log-settings module using the syntax in the following

sections.

MODIFY

modify tmm

options:

arp-log-level [debug | error | informational | notice | warning]
http-compression-log-level [debug | error | informational | notice | warning]
http-log-level [debug | error | informational | notice | warning]
ip-log-level [debug | informational | notice | warning]
irule-log-level [debug | error | informational | notice | warning]
layer4-log-level [debug | informational | notice]
net-log-level [critical | debug | error | informational | notice | warning]
os-log-level [alert | critical | debug | emergency | error | informational | notice | warning]
pva-log-level [debug | informational | notice]
ssl-log-level [alert | critical | debug | emergency | error | informational | notice | warning]

edit tmm

options:

all-properties
non-default-properties

DISPLAY

list tmm

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the tmm component to change the level of the messages about the tmm that appear in the system logs. Additionally, you can display information about the tmm.

EXAMPLES

list tmm

Displays information about the tmm.

modify tmm http-compression-log-level critical

Changes the level of the messages about HTTP compression that display in the system log to warning.

OPTIONS

arp-log-level

Specifies the lowest level of ARP messages from the tmm daemon to include in the system log. The default value is warning.

http-compression-log-level

Specifies the lowest level of HTTP compression messages from the tmm daemon to include in the system log. The default value is error.

http-log-level

Specifies the lowest level of HTTP messages from the daemon to include in the system log. The default value is error.

ip-log-level

Specifies the lowest level of IP address messages from the tmm daemon to include in the system log. The default value is warning.

irule-log-level

Specifies the lowest level of iRule messages from the tmm daemon to include in the system log. The default value is warning.

layer4-log-level

Specifies the lowest level of Layer 4 messages from the tmm daemon to include in the system log. The default value is notice.

net-log-level

Specifies the lowest level of network messages from the tmm daemon to include in the system log. The default value is warning.

os-log-level

Specifies the lowest level of operating system messages from the tmm daemon to include in the system log. The default value is notice.

pva-log-level

Specifies the lowest level of PVA messages from the tmm daemon to include in the system log. The default value is informational.

ssl-log-level

Specifies the lowest level of SSL messages from the tmm daemon to include in the system log. The default

value is warning.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-31 sys daemon-log-settings tmm(1)

sys datastor

NAME

datastor - Configures the data storage used for optimization.

MODULE

sys

SYNTAX

Configure the datastor component within the sys module using the syntax in the following sections.

MODIFY

modify datastor

options:

dedup-cache-weight [integer]

description [string]

disk [disabled | enabled]

high-water-mark [integer]

low-water-mark [integer]

web-cache-weight [integer]

DISPLAY

list datastor

show running-config datastor

options:

all-properties

cache-size

non-default-properties

one-line

store-size

DESCRIPTION

You can use the datastor component to configure disk I/O operations and optimized page cache for frequently accessed sectors. Note that symmetric data deduplication is one consumer of this storage space.

EXAMPLES

list datastor all-properties

Displays the data storage settings.

modify datastor disk disabled

Disables data storage on the disk.

OPTIONS

cache-size

Displays the size of the data storage in megabytes (MB).

dedup-cache-weight

Specifies the relative weight of the dedup cache for the Acceleration Manager module. The default value is 10.

description

User defined description.

disk Enables or disables the use of the disk (in addition to memory) for data storage.

If you enable or disable data storage on the disk, you must then restart the datastor service from the command line using the command sequence `bigstart restart datastor`.

high-water-mark

Specifies the percentage of full cache above which pruning starts. The valid range is 60 - 100 percent. The default value is 92.

low-water-mark

Specifies the percentage of full cache below which pruning stops. The valid range is 10 - 90 percent. The default value is 80.

store-size
Displays the amount of space for each disk path specified.

web-cache-weight
Specifies the relative weight of the web cache for the Acceleration Manager module. The default value is 10.

SEE ALSO

wom deduplication, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-12-06 sys datastor(1)

sys db

NAME

db - Displays or modifies bigdb database entries.

MODULE

sys

SYNTAX

Configure the db component within the sys module using the syntax in the following sections.

MODIFY

modify db [name] value [database variable value]
modify db [name] reset-to-default

DISPLAY

list db
list db [[[name] | [glob] | [regex]] ...]
options:
all-properties
default-value
non-default-properties
one-line
value
value-range

show running-config db
show running-config db [[[name] | [glob] | [regex]] ...]
options:
all-properties

DESCRIPTION

You can use the db component to modify and retrieve the data that is stored in the bigdb configuration database.

Important: After you change a bigdb database variable using the db component, you must run the command sequence save config. If you do not, the next time that you run the command sequence load [config base | config], the value of the bigdb database variable may be reset to the value in the stored configuration.

Note that tmsh only displays bigdb database entries when you explicitly request them.

EXAMPLES

modify db Connection.SynCookies.Threshold value 16384

Sets the database entry, SYN Check(tm) Activation Threshold, to the given value.

modify db Connection.SynCookies.Threshold reset-to-default

Sets the database entry, SYN Check(tm) Activation Threshold, back to the default value.

list log.mcpd.level

Displays the properties of the database entry log.mcpd.level:

OPTIONS

default-value
Displays the system-supplied default value of the database entry.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies the unique name of the database variable. This option is required for the command modify.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reset-to-default
Resets the database variable back to its default value.

value
Specifies the value to which you want to set the specified database entry.

value-range
Displays the type of data that you can use with the value option. The options are:

integer
IP address
list of valid values
management IP address
string
unsigned integer

SEE ALSO
glob, list, modify, regex, show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-08-05 sys db(1)

sys default-config

NAME
default-config - Loads the default configuration of the BIG-IP(r) system stored in the configuration files to the running configuration of the system.

MODULE
sys

SYNTAX
Configure the default-config component within the sys module using the following syntax.

MODIFY
load default-config

DESCRIPTION
You can use the default-config component to load the default system configuration to the running configuration. This results in the user-defined configuration being removed from the running configuration.

EXAMPLES
load default-config

Loads the default configuration stored on the system to the running configuration of the system.

SEE ALSO
load, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-11 sys default-config(1)

sys diags ihealth-request

NAME

ihealth-request - Displays the BIG-IP(r) iHealth qkview upload request status, errors and progress.

MODULE

sys diags

SYNTAX

Display the ihealth-request component within the sys diags module using the syntax in the following section.

DISPLAY

```
list sys diags ihealth-request
  all-properties
  one-line
```

```
show sys diags ihealth-request
  raw
```

DESCRIPTION

You can use the ihealth-request component to display the status of the iHealth upload feature.

EXAMPLES

```
list ihealth-request
```

Displays iHealth upload request status.

OPTIONS

ihealth-start-time

This is the start time of the ihealth request in seconds since 1970-01-01 00:00:00 UTC.

ihealth-finish-time

This is the finish time of the ihealth request in seconds since 1970-01-01 00:00:00 UTC.

error

This is the error status of the ihealth request: (one of none, creds, qkview-error, qkview-timeout, ihealth-unreach, ihealth-process, ihealth-sr-format, ihealth-timeout, ihealth-resolv ihealth-invalid-opts, ihealth-missing-user-password, ihealth-login-failed, ihealth-missing-qkview-file, qkview-invalid-opts, ihealth-miscellaneous-error).

status

This is the status of the ihealth upload process (one of none, running, analyzing or uploading).

progress

This is the progress in percent of the upload process.

type This is the type of the diagnostics request (one of none, ihealth, qkview or tcpdump).

qk-progress

This is the percent complete of a qkview in the process of being generated.

qk-progress-msg

Warnings and non fatal errors while running qkview

qkview-filename

This is the filename of the last qkview file generated, if there is one.

qkview-date

This is the date of the last qkview file generated, if there is one.

qkview-size

This is the file size of the last qkview file generated, if there is one.

qkview-user

This is the user name of the account that created the last qkview file generated, if there is one.

tcpdump-filename

This is the filename of the last tcpdump file generated, if there is one.

tcpdump-date

This is the date of the last tcpdump file generated, if there is one.

tcpdump-size

This is the file size of the last tcpdump file generated, if there is one.

For information about the options that you can use with the command list, see help list.

SEE ALSO

list, tmsb, qkview

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2015-2016. All rights reserved.

BIG-IP 2019-05-07 sys diags ihealth-request(1)

sys diags ihealth-result

NAME

ihealth-result - Displays the BIG-IP(r) iHealth qkview upload results.

MODULE

sys diags

SYNTAX

Display the ihealth-result component within the sys diags module using the syntax in the following section.

DISPLAY

list ihealth-result

options:

all-properties

one-line

show ihealth-result

options:

field-fmt

DESCRIPTION

You can use the ihealth-result component to display the result history of the iHealth upload feature.

EXAMPLES

list ihealth-result

Displays iHealth upload result history for the system.

OPTIONS

time-checked

This is the date and time that the qkview was uploaded to iHealth.

check-user

This is the iHealth user name (email address) used to access the iHealth service.

result

This is the result returned from iHealth.

sr This is the Support Request number uploaded to the iHealth service with the qkview file.

description

This is the description uploaded to the iHealth service with the qkview file.

size This is the size of the qkview file uploaded to the iHealth service.

id This is the iHealth ID number returned from the iHealth service for the qkview file.

warnings

This is the the total number of warnings returned by the iHealth service.

criticals

This is the the total number of critical warnings returned by the iHealth service.

For information about the options that you can use with the command list, see help list.

SEE ALSO

list, tmsh, qkview

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2015. All rights reserved.

BIG-IP 2017-01-16 sys diags ihealth-result(1)

sys diags ihealth

NAME

ihealth - Displays the BIG-IP(r) iHealth qkview upload settings.

MODULE

sys diags

SYNTAX

Display and modify the ihealth component within the sys diags module using the syntax in the following section.

MODIFY

modify ihealth

options:

user

password

expiration

options

no-ihealth

DISPLAY

list ihealth

options:

all-properties

one-line

show ihealth

options:

field-fmt

DESCRIPTION

You can use the ihealth component to display or modify the configuration of the iHealth upload feature.

EXAMPLES

list ihealth

Displays iHealth upload configuration information for the system.

modify ihealth expiration 180

Modify the expiration of iHealth upload history to 180 days.

modify ihealth options "-t 180"

Add a time limit of 180 seconds for qkview module execution.

modify ihealth no-ihealth true

Prevent any qkviews generated on this device from being uploaded to iHealth.

OPTIONS

user This is the iHealth user name (email address) used to access the iHealth service.

password

This is the iHealth password used to access the iHealth service.

expiration

This is the number of days to retain iHealth upload history. The default is 30, and the maximum is 365.

options

This is a command line string to specify qkview options.

no-ihealth

This switch causes a tag to be inserted into qkview files that instructs the iHealth service to deny their uploading.

For information about the options that you can use with the command list, see help list.

SEE ALSO

list, tmsh, qkview

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2015. All rights reserved.

BIG-IP 2018-10-04 sys diags ihealth(1)

sys disk application-volume

NAME

application-volume - Configures an application volume instance.

MODULE

sys disk

SYNTAX

Configure the application-volume component in the sys disk module using the syntax shown in the following sections.

DISPLAY

show application-volume [name]

list application-volume [name]

DELETE

delete application-volume [name]

DESCRIPTION

The application-volume component provides better granularity for managing disks. Physical disks can now be shared by several application-volumes. An application-volume is physically confined to one logical disk. The visibility of the application-volume can be confined to a particular software volume set or it can be global. No application-volume properties are allowed to be modified through tmsh or iControl(r) interfaces.

EXAMPLES

delete application-volume mysqlpdb_MD1.3

Deletes an application-volume named mysqlpdb_MD1.3.

show application-volume mysqlpdb_MD1.3

Displays the configuration details of the application-volume mysqlpdb_MD1.3 in a table.

OPTION

logical-disk [name]

Specifies the name of the logical disk in which the application-volume will be created.

owner [unassigned/datastor/mysql/vcmp]

Specifies the owner for which this application-volume is assigned. unassigned - is the default option and means the volume is not in use and nobody owns it.

preservability [discardable/precious]

Specifies the if application-volume can be discarded by software (for example, during module provisioning). discardable - is the default option.

resizeable [false/true]

Specifies the if application-volume can potentially be resized. false - is the default option.

size [integer]

Specifies the size of the application-volume.

volume-set-visibility-restraint [name]

Specifies the name of the volume set to which the application-volume is constrained, if any.

SEE ALSO

delete, show, list, tmsh, sys provision, sys disk logical-disk

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys disk application-volume(1)

sys disk directory

NAME

directory - Manages resizing of system directories.

MODULE

sys disk

SYNTAX

Configure the directory component in the sys disk module using the syntax shown in the following sections.

MODIFY
modify directory [directory_name]
options:
new-size [new_size]

SHOW
show directory

DESCRIPTION

The directory component assists in resizing system directories. It allows system administrators to increase the size of 4 system directories (/config, /shared, /var, /var/log). This allows more flexible management of the system resources and path for growing the directory sizes on case per case basis.

EXAMPLES

modify directory /shared new-size 35000

Increases the size of /shared system directory to 35 MiB.

show directory

Displays a table with currently scheduled directories for resizing. If there are no such directories the output is empty.

SEE ALSO

modify, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-11-20 sys disk directory(1)

sys disk logical-disk

NAME
logical-disk - Manages logical disks.

MODULE
sys disk

SYNTAX
Configure the logical-disk component in the sys disk module using the syntax shown in the following sections.

MODIFY
modify logical-disk [name]
options:
vg-reserved [integer]
mode [none/mixed/datastor]

DISPLAY
list logical-disk [name]

DESCRIPTION

The logical-disk component provides better granularity for managing disks. A physical disk can now be shared by one or more logical disks. A logical disk is physically confined to one physical disk.

EXAMPLES

modify logical-disk foo mode mixed vg-reserved 200

Modifies the logical disk foo mode property to mixed and the vg-reserved property size to 200 MiB.

list logical-disk foo

Displays the configuration details of the logical disk named foo.

OPTION

mode [none/mixed/datastor/control]
Specifies the current mode of the logical disk. The options are:

control - Indicates that the logical disk is part of a RAID array.

datastor - Indicates that the entire disk is committed to the datastor module.

mixed - Indicates that the disk contains multiple volumes for software and/or multiple volumes for application data.

none - Indicates that the disk is not in use. This is the default option.

size [integer]
Specifies the size (MiB) of the logical disk.

vg-free [integer]
Specifies the usable free space (MiB) available in the logical disk.

vg-in-use [integer]
Specifies the total logical disk space (MiB) in use.

vg-reserved [integer]
Specifies the reserved logical disk space (MiB). This space is NOT available for provisioning.

SEE ALSO

modify, list, tmsh, sys provision, sys disk logical-disk

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys disk logical-disk(1)

sys dns

NAME

dns - Configures the Domain Name System (DNS) for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Modify the dns component within the sys module using the syntax shown in the following sections.

MODIFY

modify dns

options:

```
description [string]
include [string]
name-servers [add | delete | replace-all-with] {
  [IP address] ...
}
name-servers none
number-of-dots [integer]
search [add | delete | replace-all-with] {
  [domain] ...
}
search none
```

edit dns

options:

```
all-properties
non-default-properties
```

DISPLAY

list dns

list dns [option]

show running-config dns

show running-config dns [option]

options:

```
all-properties
non-default-properties
one-line
```

DESCRIPTION

You can use the dns component to manage configurations by server grouping, in this case, DNS servers.

EXAMPLES

```
modify dns name-servers add { 192.168.10.20 192.168.10.22 }
```

Adds DNS name servers with the IP addresses, 192.168.10.20 and 192.168.10.22, to the BIG-IP system.

```
modify dns search add { siterequest.com store.siterequest.com london.siterequest.com }
```

Adds the host names, siterequest.com, store.siterequest.com, and london.siterequest.com, to the DNS search configuration for the BIG-IP system.

Note: When DNS searches for the host, siterequest, which is not a fully qualified domain name, it uses the IP address of the first match, in this case, siterequest.com.

```
show running-config dns
```

Displays the running configuration of the dns component.

OPTIONS

description

User defined description.

include

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using the include option. If you use this option incorrectly, you put the functionality of the system at risk.

name-servers

Configures a group of DNS name servers for the BIG-IP system.

number-of-dots

Configures the number of dots needed in a name before an initial absolute query will be made.

search

Configures a list of domain names in a specific order. DNS uses that order when searching for host names that are not fully qualified. You can use this option to delete domain names in the list.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013. All rights reserved.

BIG-IP 2015-12-07 sys dns(1)

sys dynad instrumentation

NAME

instrumentation - Display and configure the instrumentation available for the BIG-IP(r) DynaD feature.

MODULE

sys dynad

SYNTAX

Display and modify the instrumentation component within the sys dynad module using the syntax in the following section.

MODIFY

modify sys dynad instrumentation [name]

options:

active

DELETE

delete sys dynad instrumentation [name]

DISPLAY

list sys dynad instrumentation

list sys dynad instrumentation [name]

options:

all-properties

one-line

active

DESCRIPTION

You can use the instrumentation component to display or modify the configuration of the DynaD feature.

EXAMPLES

list sys dynad instrumentation

Display all the instrumentation available for use by the DynaD feature.

list sys dynad instrumentation [name]

Display a specific instrumentation available for use by the DynaD feature.

modify sys dynad instrumentation [name] active true

Activate the specified instrumentation.

delete sys dynad instrumentation [name]

Remove the specified instrumentation from the file system.

OPTIONS

active

Indicates the state of the instrumentation. The system will be instrumented when set to 'true'.

For information about the options that you can use with the command list, see help list.

SEE ALSO

list, tmsh, dynad

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2018-02-22 sys dynad instrumentation(1)

sys dynad key

NAME

key - Encryption key used for the BIG-IP(r) DynaD feature.

MODULE

sys dynad

SYNTAX

Manage the key component within the sys dynad module using the syntax in the following section.

GENERATE

generate sys dynad key

DELETE

delete sys dynad key

DISPLAY

list sys dynad key

options:

all-properties

one-line

key

blob

DESCRIPTION

You can use the key component to manage the encryption key used by the DynaD feature.

EXAMPLES

generate sys dynad key

Create a new encryption key and blob.

delete sys dynad key

Delete the existing encryption key and blob.

list sys dynad key

Display the entire key component.

OPTIONS

key Display the encryption key used to load DynaD instrumentation.

blob Display the encrypted blob used by F5 support to prepare DynaD instrumentation.

For information about the options that you can use with the command list, see help list.

SEE ALSO

list, tmsh, dynad

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-11-03 sys dynad key(1)

sys dynad rpm

NAME

rpm - Display and install instrumentation rpms for the BIG-IP(r) DynaD feature.

MODULE

sys dynad

SYNTAX

Manage instrumentation with the rpm component within the sys dynad module using the syntax in the following section.

INSTALL

install sys dynad rpm [package]

options:

force

UNINSTALL

uninstall sys dynad rpm [package]

DELETE

delete sys dynad rpm [package]

DISPLAY

list sys dynad rpm

list sys dynad rpm [package]

options:

all-properties

one-line

arch

description

path

release

version

show sys dynad rpm

show sys dynad rpm [package]

DESCRIPTION

You can use the rpm component to view and install DynaD instrumentation packages. DynaD instrumentation packages must be placed in the /shared/rpms directory to be used by the BIG-IP(r) system.

Note that installation can only occur when the DynaD feature is inactive.

EXAMPLES

list sys dynad rpm

Display all the rpms available for installation.

list sys dynad rpm [package]

Display a specific rpm available for installation.

install sys dynad rpm [package]

Install the specified instrumentation rpm.

install sys dynad rpm force [package]

Forcibly install the specified instrumentation rpm.

uninstall sys dynad rpm [package]

Remove the specified instrumentation package from the rpm database.

delete sys dyad rpm [package]

Delete the specified instrumentation package from the BIG-IP(r) file system.

show sys dynad rpm

Display all previously installed instrumentation rpms.

OPTIONS

force

Forcibly install the specified package regardless of most errors. This option is useful if an RPM needs to be reinstalled.

arch Displays the architecture of the instrumentation contained in the package.

description
Displays the description of the instrumentation contained in the package.

path Displays the local path to the package.

release
Displays the release number of the package.

version
Displays the version number of the package.

For information about the options that you can use with the command list, see help list.

SEE ALSO
list, tmsh, dynad

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2018-02-22 sys dynad rpm(1)

sys dynad settings

NAME
Settings - Display and configure global settings for the BIG-IP(r) DynaD feature.

MODULE
sys dynad

SYNTAX
Display and modify the settings component within the sys dynad module using the syntax in the following section.

MODIFY
modify sys dynad settings
options:
development-mode

DISPLAY
list sys dynad settings
options:
all-properties
one-line
development-mode

DESCRIPTION
You can use the settings component to display or modify the global settings for the DynaD feature.

EXAMPLES
list sys dynad settings

Display all the settings available to the DynaD feature.

modify sys dynad settings development-mode true

Enable development mode.

OPTIONS
development-mode
Development-mode can be used to place the DynaD feature into a promiscuous operating mode. Allowing it to bypass various validation checks performed on installed scripts.

For information about the options that you can use with the command list, see help list.

SEE ALSO
list, tmsh, dynad

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016-2017. All rights reserved.

BIG-IP 2017-09-05 sys dynad settings(1)

sys dynad status

NAME

status - Displays the status of the BIG-IP(r) DynaD feature.

MODULE

sys dynad

SYNTAX

Display the status component within the sys dynad module using the syntax in the following section.

DISPLAY

show sys dynad status

DESCRIPTION

You can use the status component to display the overall state of the DynaD feature.

EXAMPLES

show sys dynad status

Displays status of the DynaD feature.

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh, dynad

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2016. All rights reserved.

BIG-IP 2017-05-19 sys dynad status(1)

sys ecm config

NAME

config - Configures Elastic Compute Manager cluster.

MODULE

sys ecm

SYNTAX

Modify the configuration for ECM module using the syntax shown in the following sections.

MODIFY

modify config

options:

seed-ip [string]

dns-resolver [string]

edit config

options:

all-properties

non-default-properties

DISPLAY

list config

options:

auth [string]

seed-ip [string]

dns-resolver [string]

status [string]

OPTIONS

You can use these options with the config component:

auth Displays the authorization credentials to access ECM cluster. Authorization credentials are obtained by registering the BIG-IP with the ECM cluster. Refer to register command within the ECM module.

seed-ip
Specifies the host name or IP address of the seed node of the ECM cluster.

dns-resolver
Specifies the BIG-IP DNS resolver to use for host name resolution.

status
Displays the status

SEE ALSO
edit, list, modify, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-10-26 sys ecm config(1)

sys ecm register

NAME
register - Registers this BIG-IP with ECM cluster.

MODULE
sys ecm

SYNTAX
Initiate the registration of this BIG-IP with the ECM cluster. Registration enables the BIG-IP to access services running in the ECM cluster.

RUN
run register

SEE ALSO
edit, list, modify, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2017-05-24 sys ecm register(1)

sys failover

NAME
failover - Configures failover for a BIG-IP(r) unit in a redundant system configuration.

MODULE
sys

SYNTAX
Change the failover state within the sys module using the syntax in the following section.

MODIFY
run failover
options:
device [string]
no-persist
offline
online
persist
standby
traffic-group [[string] | default | non-default | none]

DISPLAY
show failover

options:
cable

DESCRIPTION

Failover is the process where a standby unit in a redundant system configuration takes over when a software or hardware failure is detected on the active unit.

EXAMPLES

run failover standby

Causes the active unit or cluster to go into the standby state forcing the other unit or cluster in the redundant system configuration to become active.

run failover offline

Causes the active unit or cluster to go into the Forced Offline state.

run failover online

Changes the status of a unit or cluster from Forced Offline to either Active or Standby, depending upon the status of the other unit or cluster in a redundant system configuration.

show failover

Displays the failover state of the BIG-IP system (active, standby, offline, forced_offline) and how long it has been in that state.

run failover standby device my_bigip

Specifies that the my_bigip device should become the active device for all traffic groups currently active on this device.

run failover standby traffic-group traffic_grp01

Specifies that the traffic group named traffic_grp01 should fail over to the Standby state. The traffic group will then become Active on another device.

run sys failover offline no-persist

Changes the status of a unit to Forced Offline and indicates that the change will not be persisted after a system restart.

run sys failover offline persist

Changes the status of a unit to Forced Offline and indicates that the change will be persisted after a system restart.

OPTIONS

Use these options to control failover of the system:

device

Specifies the device that should next become the active device for the specified traffic group or all traffic groups (if a traffic group is not specified). This option may only be specified with the standby option.

no-persist

Does not persist the change in status of a unit. The option is valid only with the offline state.

offline

Changes the status of a unit or cluster to Forced Offline. If persist or no-persist options are not specified, the default action is to persist the offline status of the unit between system restarts.

online

Changes the status of a unit or cluster from Forced Offline to either Active or Standby, depending upon the status of the other unit or cluster in a redundant system configuration.

persist

Persists the change in status of a unit. The option is valid only with the offline state.

standby

Specifies that the active unit or cluster fails over to a Standby state, causing the standby unit or cluster to become Active.

traffic-group

Specifies the traffic-group that should fail over to the Standby state, the traffic-group will become Active on another device. This option may only be specified with the standby option.

Use this option to display the failover cable status of the system:

cable

Displays the status that the failover daemon detects on the serial cable from its failover peer. It also shows what the failover peer detects on the serial cable. An active BIG-IP system will see a zero from its failover peer. A standby BIG-IP system will see a one from its failover peer.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2017-03-13 sys failover(1)

sys feature-module

NAME

feature-module - Enables or disables a feature module on the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the feature-module component within the sys module using the syntax in the following sections.

MODIFY

modify feature-module
modify feature-module [[all]]
options:
enabled | disabled

edit feature-module
[[glob] | [regex]] ...]
options:
all-properties
non-default-properties

DISPLAY

list feature-module
list feature-module
[[glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line

DESCRIPTION

You can use the feature-module component to modify the availability of any licensed feature modules on your system.

EXAMPLES

list feature-module
Displays the current feature module of the system.

OPTIONS

all Specifies that you are enabling or disabling all of the available modules.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

edit, glob, list, modify, regex, show, tmsh, provision

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013. All rights reserved.

BIG-IP 2019-02-21 sys feature-module(1)

sys file apache-ssl-cert

NAME

apache-ssl-cert - Manages an Apache SSL certificate file.

MODULE

sys file

SYNTAX

Configure the apache-ssl-cert component within the sys file module using the syntax shown in the following sections.

CREATE

create apache-ssl-cert [name]

options:

source-path [URL]

DISPLAY

list apache-ssl-cert

list apache-ssl-cert [[[name] | [glob] | [regex]] ...]

DELETE

delete apache-ssl-cert [name]

DESCRIPTION

You can use the apache-ssl-cert component to create, delete, or list an SSL certificate.

EXAMPLES

create apache-ssl-cert new-cert source-path http://cert-server/cert_store/certs/cert1.crt

Downloads the certificate from the given URL into file-store, creates an SSL certificate file named new-cert, and saves the given URL in the source-path attribute.

create apache-ssl-cert new-cert source-path file:/shared/save/cert1.crt

Specifies the location of the file on the local disk (use this when the file has already been created on the local disk).

SUPPORTED URL FORMAT

Supported URL schemes are HTTP, HTTPS, FTP, FTPS, and FILE.

OPTIONS

bundle-certificates

Lists data about all the certificates in the bundle, if the certificate file is a bundle. This option must be explicitly specified; otherwise, this field will be none.

certificate-key-curve-name

Specifies the Elliptical Curve name of the cryptographic key associated with this certificate. This field will be set to none if an Elliptical Curve key is not present.

certificate-key-size

Specifies the number of bits in the key associated with this certificate.

checksum

Specifies a cryptographic hash or checksum of the file contents for use in verification of file integrity.

create-time

Specifies the time at which the file-object was created.

created-by

Specifies the user who originally created the file-object.

expiration-date

Specifies the date at which this certificate expires. Stored as a POSIX time.

expiration-string

Specifies a string representation of the expiration date of the certificate.

fingerprint

Displays the SHA-256 fingerprint of the certificate.

is-bundle

Specifies whether the certificate file is a bundle (that is, whether it contains more than one certificate).

issuer

Specifies X509 information of the certificate's issuer. If the cert is a bundle, this displays the issuer information for the primary (first) cert in the bundle.

key-type

Specifies the type of cryptographic key associated with this certificate.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

partition
Specifies the administrative partition where the certificate resides.

revision
Identifies the latest revision of the file. The revision starts with 1 and gets incremented on each update.

serial-number
Specifies the certificate's serial number.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]
This attribute takes a URL, for example:

source-path http://cert-server/cert_store/certs/vs_132.crt

source-path https://cert-server/cert_store/certs/vs_132.crt

source-path ftp://username:password@server/cert_store/certs/vs_132.crt

subject
Specifies X509 information of the certificate's subject. If the cert is a bundle, this displays the subject information for the primary (first) cert in the bundle.

subject-alternative-name
Specifies a standard X.509 extension as shown in RFC 2459.

updated-by
Specifies the user who last updated the file-object.

version
Specifies the X509 version of the certificate.

SEE ALSO

create, delete, glob, list, ltm profile client-ssl, ltm profile server-ssl, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2017-11-16 sys file apache-ssl-cert(1)

sys file browser-capabilities-db

NAME

browser-capabilities-db - Manages a browser capabilities DB file.

MODULE

sys file

SYNTAX

Configure the browser-capabilities-db component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

create browser-capabilities-db [name]

modify browser-capabilities-db [name]

options:

source-path [URL]

DISPLAY

list browser-capabilities-db

list browser-capabilities-db [[[name] | [glob] | [regex]] ...]

DELETE

delete browser-capabilities-db [name]

DESCRIPTION

You can use the browser-capabilities-db component to create, delete, list or modify an browser capabilities DB file.

EXAMPLES

create browser-capabilities-db dcdb source-path file:/shared/images/dcdb

Loads the browser capabilities file from the given path on the local disk into file-store and creates an file named dcdb.

OPTIONS

checksum

A cryptographic hash or checksum of the file contents for use in verification of file integrity.

create-time

Specifies the time at which the file-object was created.

created-by

Specifies the user who originally created the file-object.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision

Specifies the latest revision of the file. The revision starts with 1 and gets incremented on each update.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [path]

This attribute takes an absolute path on the local disk, for example:

source-path file:/shared/images/filename

updated-by

Specifies the user who last updated the file-object.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-12-29 sys file browser-capabilities-db(1)

sys file data-group

NAME

data-group - Manages an external data group file.

MODULE

sys file

SYNTAX

Manage the data-group component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

create data-group [name]

modify data-group [name]

options:

app-service [[string] | none]

data-group-description [string]

data-group-name [name]

separator [string]

source-path [URL]

type [integer | ip | string]

edit data-group [[[name] | [glob] | [regex]] ...]

DISPLAY

list data-group

list data-group [[[name] | [glob] | [regex]] ...]

DELETE

delete data-group [name]

DESCRIPTION

You can use the data-group component to create, edit, delete, list or modify an external data group file.

EXAMPLES

```
create data-group new-dg source-path http://file-server/data-groups/acl.class type string
```

Downloads the data-group file from the given URL into file-store, creates an external-data-group file named new-dg, and saves the given URL in the source-path attribute.

```
create data-group new-dg source-path http://file-server/data-groups/acl.class type string data-group-name dg data-group-description "created for rule xyz"
```

Downloads the data-group file from the given URL into file-store, creates an external-data-group file named new-dg, saves the given URL in the source-path attribute, and creates an external data group within the Itm data-group module named dg with the given description.

```
create data-group new-dg source-path file://shared/save/Test.cls type ip
```

Specifies the location of the file on the local disk (use this when the file has already been created on the local disk).

SUPPORTED URL FORMAT

Supported URL schemes are HTTP, HTTPS, FTP, FTPS, and FILE.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum

Specifies a cryptographic hash or checksum of the file contents for use in verification of file integrity.

created-by

Specifies the user who originally created the file-object.

create-time

Specifies the time at which the file-object was created.

data-group-description

Specifies the description of the external data group that will be created within the Itm data-group module and reference the given data group file. This is optional in the create command.

data-group-name

Specifies the name of the external data group that will be created within the Itm data-group module and reference the given data group file. This is optional in the create command.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision

The latest revision of the file. The revision starts with 1 and gets incremented on each update.

separator

Specifies a separator to use when defining the data group. The default value is :=.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]

This attribute takes a URL, for example:

```
source-path http://file-server/data-groups/AUL_1.cls
```

```
source-path https://file-server/data-groups/CNN.x
```

```
source-path ftp://username:password@server/data-groups/latest.class
```

```
source-path file://shared/save/Test.dat
```

type Specifies the kind of data in the group. This option is required by the create command.

Possible values for type are:

• integer

• ip

• string

updated-by

Specifies the user who last updated the file-object.

SEE ALSO

create, delete, edit, glob, list, Itm data-group external, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-05-22 sys file data-group(1)

sys file device-capabilities-db

NAME

device-capabilities-db - Manages a device capabilities DB file.

MODULE

sys file

SYNTAX

Configure the device-capabilities-db component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

```
create device-capabilities-db [name]
modify device-capabilities-db [name]
options:
  app-service [[string] | none]
  source-path [file:/PATH/FILE]
```

DISPLAY

```
list device-capabilities-db
list device-capabilities-db [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete device-capabilities-db [name]
```

DESCRIPTION

You can use the device-capabilities-db component to create, delete, list or modify an device capabilities DB file.

EXAMPLES

```
create device-capabilities-db dcdb source-path file:/shared/images/dcdb
```

Loads the device capabilities file from the given path on the local disk into file-store and creates an file named dcdb.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum

A cryptographic hash or checksum of the file contents for use in verification of file integrity.

create-time

Specifies the time at which the file-object was created.

created-by

Specifies the user who originally created the file-object.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision

Specifies the latest revision of the file. The revision starts with 1 and gets incremented on each update.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [path]

This attribute takes an absolute path on the local disk, for example:

```
source-path file:/shared/images/filename
```

updated-by

Specifies the user who last updated the file-object.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2018-02-07 sys file device-capabilities-db(1)

sys file external-monitor

NAME

external-monitor - Manages an external monitor file.

MODULE

sys file

SYNTAX

Manage the external-monitor component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

```
create external-monitor [name]
modify external-monitor [name]
options:
  app-service [[string] | none]
  source-path [URL]
```

```
edit external-monitor [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list external-monitor
list external-monitor [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete external-monitor [name]
```

DESCRIPTION

You can use the external-monitor component to create, edit, delete, list or modify an external-monitor file.

EXAMPLES

```
create external-monitor new-mon source-path http://file-server/external-monitors/mon_app1
```

Downloads the monitor file from the given URL into file-store, creates an external-monitor file named new-mon, and saves the given URL in the source-path attribute.

```
create external-monitor new-mon source-path file://shared/save/Test.mon
```

Specifies the location of the file on the local disk (use this when the file has already been created on the local disk).

SUPPORTED URL FORMAT

Supported URL schemes are HTTP, HTTPS, FTP, FTPS, and FILE.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum

Specifies a cryptographic hash or checksum of the file contents for use in verification of file integrity.

created-by

Specifies the user who originally created the file-object.

create-time

Specifies the time at which the file-object was created.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical

value.

revision
The latest revision of the file. The revision starts with 1 and gets incremented on each update.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]
This attribute takes a URL, for example:

source-path http://file-server/external-monitors/monitor_service

source-path https://file-server/external-monitors/custom_mon.1

source-path ftp://username:password@server/external-monitors/tested.mon

updated-by
Specifies the user who last updated the file-object.

SEE ALSO

create, delete, edit, glob, list, ltm monitor external, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-10-19 sys file external-monitor(1)

sys file ifile

NAME
ifile - Manages an iFile file.

MODULE
sys file

SYNTAX
Manage the ifile component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY
create ifile [name]
modify ifile [name]
options:
app-service [[string] | none]
source-path [URL]

edit ifile [[[name] | [glob] | [regex]] ...]

DISPLAY
list ifile
list ifile [[[name] | [glob] | [regex]] ...]

DELETE
delete ifile [name]

DESCRIPTION
You can use the ifile component to create, edit, delete, list or modify an iFile file.

EXAMPLES
create ifile new-ifile source-path http://tmp/text.txt

Downloads the iFile file from the given URL into file-store and creates an ifile file named new-ifile. Saves the given URL in the source-path attribute.

Supported URL schemes are "HTTP", "HTTPS", "FTP", "FTPS" & "FILE"

OPTIONS
app-service
Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum
A cryptographic hash or checksum of the file contents for use in verification of file integrity.

created-by

Specifies the user who originally created the file-object.

create-time
Specifies the time at which the file-object was created.

last-update-time
Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision
The latest revision of the file. The revision starts with 1 and gets incremented on each update.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]
This attribute takes a URL, for example:

source-path http://file-server/ifiles/AUL_1.cls

source-path https://file-server/ifiles/CNN.x

source-path ftp://username:password@server/ifiles/latest.class

updated-by
Specifies the user who last updated the file-object.

SEE ALSO

create, delete, edit, glob, list, ltm ifile, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 sys file ifile(1)

sys file lwtunneltbl

NAME
lwtunneltbl - Manages an lwtunneltbl file.

MODULE
sys file

SYNTAX
Manage the lwtunneltbl component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY
create lwtunneltbl [name]
modify lwtunneltbl [name]
options:
app-service [[string] | none]
source-path [URL]

edit lwtunneltbl [[name] | [glob] | [regex]] ...]

DISPLAY
list lwtunneltbl
list lwtunneltbl [[name] | [glob] | [regex]] ...]

DELETE
delete lwtunneltbl [name]

DESCRIPTION
You can use the lwtunneltbl component to create, edit, delete, list or modify an lwtunneltbl file.

EXAMPLES
create lwtunneltbl new-ifile source-path http://tmp/text.txt

Downloads the lwtunneltbl file from the given URL into file-store and creates an lwtunneltbl file named new-ifile. Saves the given URL in the source-path attribute.

Supported URL schemes are "HTTP", "HTTPS", "FTP", "FTPS" & "FILE"

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum

A cryptographic hash or checksum of the file contents for use in verification of file integrity.

created-by

Specifies the user who originally created the file-object.

create-time

Specifies the time at which the file-object was created.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision

The latest revision of the file. The revision starts with 1 and gets incremented on each update.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]

This attribute takes a URL, for example:

source-path http://file-server/ifiles/AUL_1.cls

source-path https://file-server/ifiles/CNN.x

source-path ftp://username:password@server/ifiles/latest.class

updated-by

Specifies the user who last updated the file-object.

SEE ALSO

create, delete, edit, glob, list, ltm ifile, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2015-10-01 sys file lwtunneltbl(1)

sys file rewrite-rule

NAME

rewrite-rule - Manages a HTML content rewrite rule.

MODULE

sys file

SYNTAX

Configure the rewrite-rule component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

create rewrite-rule [name]

modify rewrite-rule [name]

options:

local-path [URL]

edit rewrite-rule [[[name] | [glob] | [regex]] ...]

DISPLAY

list rewrite-rule

list rewrite-rule [[[name] | [glob] | [regex]] ...]

DELETE

delete rewrite-rule [name]

DESCRIPTION

You can use the rewrite-rule component to create, edit, delete, list or modify a HTML content rewrite rule.

EXAMPLES

```
create rewrite-rule new-rule local-path /shared/tmp/my_rewrite_rule
```

Creates a new HTML content rewrite rule using file located by local-path and saves path in the local-path attribute.

OPTIONS

checksum

Specifies a cryptographic hash or checksum of the file contents for use in verification of file integrity.

created-by

Specifies the user who originally created the file-object.

create-time

Specifies the time at which the file-object was created.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision

Specifies the latest revision of the file. The revision starts with 1 and gets incremented on each update.

size Specifies the size (in bytes) of the file associated with this file object.

local-path [path]

This attribute takes a path, for example:

```
local-path /shared/tmp/my_rewrite_rule
```

updated-by

Specifies the user who last updated the file-object.

SEE ALSO

create, delete, edit, glob, list, ltm profile html, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 sys file rewrite-rule(1)

sys file ssl-cert

NAME

ssl-cert - Manages a SSL certificate file.

MODULE

sys file

SYNTAX

Configure the ssl-cert component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ssl-cert [name]
```

options:

```
app-service [[string] | none]
```

```
source-path [URL]
```

```
modify ssl-cert [name]
```

options:

```
app-service [[string] | none]
```

```
cert-validation-options [none | ocsp]
```

```
cert-validators [none | [cert_validator_name]]
```

```
issuer-cert [none | [issuer_cert_name]]
```

```
source-path [URL]
```

```
edit ssl-cert [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list ssl-cert
```

```
list ssl-cert [ [name] | [glob] | [regex] ] ... ]
```

DELETE

delete ssl-cert [name]

DESCRIPTION

You can use the ssl-cert component to create, edit, delete, list or modify an SSL certificate.

EXAMPLES

create ssl-cert new-cert source-path http://cert-server/cert_store/certs/cert1.crt

Downloads the certificate from the given URL into file-store, creates an SSL certificate file named new-cert, and saves the given URL in the source-path attribute.

create ssl-cert new-cert source-path file://shared/save/cert1.crt

Specifies the location of the file on the local disk (use this when the file has already been created on the local disk).

SUPPORTED URL FORMAT

Supported URL schemes are HTTP, HTTPS, FTP, FTPS, and FILE.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

bundle-certificates

Lists data about all the certificates in the bundle, if the certificate file is a bundle. This option must be explicitly specified; otherwise, this field will be none.

cert-validation-options

Specifies the option used for validating the certificate status.

cert-validators

Specifies the name of the cert-validators used for validating the certificate status. At most one cert-validator can be configured for each cert-validation type.

certificate-key-size

Specifies the number of bits in the key associated with this certificate.

checksum

Specifies a cryptographic hash or checksum of the file contents for use in verification of file integrity.

create-time

Specifies the time at which the file-object was created.

created-by

Specifies the user who originally created the file-object.

expiration-date

Specifies the date at which this certificate expires. Stored as a POSIX time.

expiration-string

Specifies a string representation of the expiration date of the certificate.

fingerprint

Displays the SHA-256 fingerprint of the certificate.

is-bundle

Specifies whether the certificate file is a bundle (that is, whether it contains more than one certificate).

issuer

Specifies X509 information of the certificate's issuer. If the cert is a bundle, this displays the issuer information for the primary (first) cert in the bundle.

issuer-cert

Specifies the name of the issuer certificate for this certificate.

key-type

Specifies the type of cryptographic key associated with this certificate.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision

Specifies the latest revision of the file. The revision starts with 1 and gets incremented on each update.

serial-number

Specifies the certificate's serial number.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]

This attribute takes a URL, for example:

source-path http://cert-server/cert_store/certs/vs_132.crt

source-path https://cert-server/cert_store/certs/vs_132.crt

source-path ftp://username:password@server/cert_store/certs/vs_132.crt

subject

Specifies X509 information of the certificate's subject. If the cert is a bundle, this displays the subject information for the primary (first) cert in the bundle.

subject-alternative-name

Specifies a standard X.509 extension as shown in RFC 2459.

updated-by

Specifies the user who last updated the file-object.

version

Specifies the X509 version of the certificate.

SEE ALSO

create, delete, edit, glob, list, ltm profile client-ssl, ltm profile server-ssl, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2017-05-05 sys file ssl-cert(1)

sys file ssl-crl

NAME

ssl-crl - Manages a SSL CRL file.

MODULE

sys file

SYNTAX

Configure the ssl-crl component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

create ssl-crl [name]

modify ssl-crl [name]

options:

app-service [[string] | none]

source-path [URL]

edit ssl-crl [[[name] | [glob] | [regex]] ...]

DISPLAY

list ssl-crl

list ssl-crl [[[name] | [glob] | [regex]] ...]

DELETE

delete ssl-crl [name]

DESCRIPTION

You can use the ssl-crl component to create, edit, delete, list or modify an SSL CRL file.

EXAMPLES

create ssl-crl new-crl source-path http://cert-server/cert_store/CRLs/latest.crl

Downloads the CRL file from the given URL into file-store, creates an SSL CRL file named new-crl, and saves the given URL in the source-path attribute.

create ssl-crl new-crl source-path file://shared/save/copy_10.crl

Specifies the location of the file on the local disk (use this when the file has already been created on the local disk).

SUPPORTED URL FORMAT

Supported URL schemes are HTTP, HTTPS, FTP, FTPS, and FILE.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum

Specifies a cryptographic hash or checksum of the file contents for use in verification of file integrity.

created-by

Specifies the user who originally created the file-object.

create-time

Specifies the time at which the file-object was created.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

revision

Specifies the latest revision of the file. The revision starts with 1 and gets incremented on each update.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]

This attribute takes a URL, for example:

source-path http://cert-server/cert_store/CRLs/backup_10.crl

source-path https://cert-server/cert_store/CRLs/jan_2010.crl

source-path ftp://username:password@server/cert_store/CRLs/latest.crl

updated-by

Specifies the user who last updated the file-object.

SEE ALSO

create, delete, edit, glob, list, ltm profile client-ssl, ltm profile server-ssl, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-04-12 sys file ssl-crl(1)

sys file ssl-key

NAME

ssl-key - Manages a SSL certificate key file.

MODULE

sys file

SYNTAX

Configure the ssl-key component within the sys file module using the syntax shown in the following sections.

CREATE/MODIFY

create ssl-key [name]

modify ssl-key [name]

options:

app-service [[string] | none]

source-path [URL]

passphrase [passphrase]

edit ssl-key [[[name] | [glob] | [regex]] ...]

DISPLAY

list ssl-key

list ssl-key [[[name] | [glob] | [regex]] ...]

DELETE

delete ssl-key [name]

DESCRIPTION

You can use the ssl-key component to create, edit, delete, list or modify an SSL certificate key file.

EXAMPLES

create ssl-key new-key source-path http://cert-server/cert_store/certs/cert1.key

Downloads the certificate-key file from the given URL into file-store and creates an SSL certificate key file named new-key. Saves the given URL in the source-path attribute.

create ssl-key new-key source-path file:/shared/save/cert1.key

Specifies the location of the file on the local disk. Use this when the file has already been created on the local disk.

SUPPORTED URL FORMAT

Supported URL schemes are HTTP, HTTPS, FTP, FTPS, and FILE.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

checksum

A cryptographic hash or checksum of the file contents for use in verification of file integrity.

create-time

Specifies the time at which the file-object was created.

created-by

Specifies the user who originally created the file-object.

key-size

Specifies the size of the cryptographic key associated with this file object, in bits.

key-type

Specifies the cryptographic type of the key in question. That is, which algorithm this key is compatible with.

The options are:

rsa-private

The key is an RSA private key.

dsa-private

The key is a DSA based private key.

last-update-time

Specifies the last time at which the file-object was updated/modified.

mode Specifies the UNIX file permissions mode for the file associated with this file-object as a numerical value.

passphrase [passphrase]

Specifies an optional passphrase with which the key has been protected. It may be used by consumers of the key in the data-plane or control-plane to decrypt it.

revision

Specifies the latest revision of the file. The revision starts with 1 and gets incremented on each update.

security-type

Specifies the type of security used to handle or store the key.

The options are:

normal

The key resides in a standard form on the file-system. This is the default value.

fips The key is protected by a FIPS device on the system and is only applicable to devices with FIPS support.

password

Specifies that the key is protected by a passphrase and stored in encrypted form.

nethsm

The key is protected by a FIPS device outside the system.

size Specifies the size (in bytes) of the file associated with this file object.

source-path [URL]

This attribute takes a URL, for example:

source-path http://cert-server/cert_store/certs/vs_132.key

source-path https://cert-server/cert_store/certs/vs_132.key

source-path ftp://username:password@server/cert_store/certs/vs_132.key

updated-by
Specifies the user who last updated the file-object.

SEE ALSO

create, delete, edit, glob, list, ltm profile client-ssl, ltm profile server-ssl, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2015. All rights reserved.

BIG-IP 2015-07-22 sys file ssl-key(1)

sys fipsuser

NAME

FIPS - Manage FIPS User Table

MODULE

sys fipsuser

SYNTAX

Manage FIPS user table with the sys fipsuser module using the syntax in the following section.

LIST

list sys fipsuser [fips-user-id]

DELETE

delete sys fipsuser {fips-user-id| all}

DESCRIPTION

You can use the fipsuser component to list or delete the one entry comprising the FIPS Cryptographic Username and encrypted password, contained in table fipsuser.

EXAMPLES

list sys fipsuser

Displays the list of the FIPS Cryptographic Username and password stored in the FIPS User Table fipsuser. The password is shown encrypted and base-64 encoded.

delete sys fipsuser f5cu

Deletes the FIPS entry for the FIPS Cryptographic Username f5cu. This removes the Username and encrypted password.

delete sys fipsuser all

Deletes all FIPS Cryptographic Usernames and corresponding passwords from the table fipsuser. This empties the table.

OPTIONS

None

SEE ALSO

list, delete, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2020. All rights reserved.

BIG-IP 2020-01-29 sys fipsuser(1)

sys fix-connection

NAME

fix-connection - Displays FIX connection statistics.

MODULE

sys

SYNTAX**DISPLAY**

show fix-connection

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

all-properties

save-to-file

DESCRIPTION

You can use the fix-connection component to display statistics about FIX connections.

EXAMPLES

show fix-connection

Displays FIX connection statistics in the system default units.

OPTIONS

For information about the options that you can use with the command show, see help show.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2014. All rights reserved.

BIG-IP 2014-05-19 sys fix-connection(1)

sys folder

NAME

folder - Configure folders (directory structure) on the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the folder component within the sys module using the syntax in the following sections.

CREATE/MODIFY

create folder [name]

modify folder [name]

options:

app-service [[string] | none]

description [string]

device-group [[string] | default | non-default | none]

no-ref-check [false | true]

traffic-group [[string] | default | non-default | none]

DISPLAY

list folder

list folder [[name] | [glob] | [regex] | [recursive]]

DELETE

delete folder [name]

DESCRIPTION

The folder system enables users to create logical containers for the purpose of granular control of synchronization to other devices in a device group.

The folder system is hierarchical, with folders and sub-folders, in a parent-to-child relationship. The highest level folder in the system is called root. For every administrative partition on the BIG-IP system, there is a top-level folder. Top-level folders always have root as the parent. Users can create sub-folders to any folder in the system.

EXAMPLES

create sys folder sub-folder1 device-group dg1 traffic-group none

Creates a new sub-folder to the current working folder called sub-folder1, associates the folder with a device-group called dg1, and sets the traffic-group to no association.

modify sys folder /Common/sub-folder1/subfolder2 description "store pools for the B2 server configuration"

Changes the description property of the folder indicated by its full name.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

device-group

Adds this folder and all configuration items in this folder to a device group for device failover or config-sync purposes. The options are:

default

Indicates that this folder should use the device group setting of its parent folder. If the parent folder's associated device group is changed, this folder's device group will change as well.

non-default

Disassociates this folder from its parent folder's device group setting. This folder's device group field can then be set independently of the parent folder's field.

hidden

Folders may be hidden by setting this property to true. The -hidden command-line option will allow you to view hidden folders, but is not required to use or modify a folder. The -hidden command-line option only affects output from the list command and the results of tab completing a configuration item. If set to false, the folder will always be visible as long as the user has the appropriate permissions.

inherited-devicegroup

Specifies, when set to true, that this folder uses the device group setting of its parent folder. If the parent folder's associated device group is changed then this folder's device group will change as well. This field is read-only.

inherited-traffic-group

Specifies, when set to true, that this folder uses the traffic group setting of its parent folder. If the parent folder's associated traffic group is changed then this folder's traffic group will change as well. This field is read-only.

no-ref-check

Specifies whether strict device group reference validation is performed on configuration items in the folder. The options are:

false

Requires configuration items in the folder to sync to a super-set of the devices that are associated with any configuration that refers to configuration items in the folder. This is the default value.

true Disables this check. It is then assumed that any dependent configuration items contained in the folder will be created locally on the other devices.

traffic-group

Adds this folder and its configuration items to an existing traffic group. The values default and non-default work as they do for the device-group option.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, include photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 sys folder(1)

sys fpga firmware-config

NAME

firmware-config - Configures the FPGA firmware to be used by the system.

MODULE

sys fpga

SYNTAX

Configure the firmware-config component within the sys fpga module using the syntax shown in the following sections.

MODIFY

modify firmware-config

options:

type [l4-performance-fpga | l7-intelligent-fpga | standard-balanced-fpga | traffic-acceleration-fpga]

DISPLAY

list firmware-config

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the firmware-config component to configure the PFGA firmware type to use.

EXAMPLES

list firmware-config

Displays properties of the current FPGA firmware configuration.

modify firmware-config type

Modify the type of the current FPGA firmware configuration. The default is standard-balanced-fpga.

OPTIONS

type The type for FPGA firmware current used on the system.

l4-performance-fpga: High throughput fpga firmware. **l7-intelligent-fpga:** eFAD and L7 intelligent fpga firmware. **standard-balanced-fpga:** The balanced standard fpga firmware. **traffic-acceleration-fpga:** traffic acceleration fpga firmware.

SEE ALSO

list, modify, tmsh, fpga

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2014. All rights reserved.

BIG-IP 2016-10-10 **sys fpga firmware-config(1)**

sys fpga info

NAME

info - Displays current FPGA (Field-Programmable Gate Array) firmware information on the system.

MODULE

sys fpga

SYNTAX

Displays current info component within the sys fpga module using the syntax in the following section.

DISPLAY

show info

options:

all-properties
field-fmt

DESCRIPTION

You can use the info component to display the current FPGA firmware information on the system.

EXAMPLES

show info

Displays current FPGA (Field-Programmable Gate Array) firmware information on the system

SEE ALSO

show, tmsh, fpga

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-03-26 **sys fpga info(1)**

sys fpga turboflex-profile

NAME

turboflex-profile - Displays current FPGA (Field-Programmable Gate Array) firmware and the Turboflex profiles it supports.

MODULE

sys fpga

SYNTAX

Displays current turboflex-profile component within the sys fpga module using the syntax in the following section.

DISPLAY

show turboflex-profile

options:

all-properties

field-fmt

DESCRIPTION

You can use the turboflex-profile component to display the current FPGA firmware & Turboflex profile on the system.

EXAMPLES

show turboflex-profile

Displays current FPGA (Field-Programmable Gate Array) firmware and the Turboflex profile it supports on the system

SEE ALSO

show, tmsh, fpga

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014, 2017. All rights reserved.

BIG-IP 2017-09-05 sys fpga turboflex-profile(1)

sys geoip

NAME

geoip - Loads the GeoIP data files.

MODULE

sys

SYNTAX

Use the geoip component within the gtm module to load the GeoIP data files using the syntax in the following sections.

LOADING

load geoip

DESCRIPTION

The BIG-IP system ships with three default database files that are stored in the `/usr/share/GeoIP/v2` directory. The three files are: `F5GeoIP.dat`, `F5GeoIPISP.dat`, and `F5GeoIPv6.dat`.

You can download and install updated GeoIP database files using the procedure available from the F5 download site. The installation places the updated database files in the `share/GeoIP/v2` directory.

When you run the `load geoip` command sequence, the system loads the GeoIP files from disk into the running configuration. If you have downloaded and installed updated database files, those files are loaded from the `/shared/GeoIP/v2` directory. Otherwise, the default database files are loaded from the `/usr/share/GeoIP/v2` directory. Note that if both directories contain the same files, the files in `shared/GeoIP/v2` are loaded.

EXAMPLES

load geoip

Loads the GeoIP files from disk into the running configuration.

SEE ALSO

load, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2015-09-21 sys geoip(1)

sys global-settings

NAME

global-settings - Configures the global system settings for a BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the global-settings component within the sys module using the syntax in the following sections.

MODIFY

modify global-settings

options:

aws-access-key [string]
aws-secret-key [string]
aws-api-max-concurrency [integer]
file-blacklist-path-prefix [string]
file-blacklist-read-only-path-prefix [string]
file-whitelist-path-prefix [string]
console-inactivity-timeout [integer]
custom-addr [IP address]
description [string]
failsafe-action [go-offline | reboot | restart-all |
go-offline-restart-tm | failover-restart-tm]
file-local-path-prefix [local path prefix]
gui-audit [disabled | enabled]
gui-expired-cert-alert [disabled | enabled]
gui-security-banner [disabled | enabled]
gui-security-banner-text [string]
gui-setup [disabled | enabled]
host-addr-mode [custom | management | state-mirror]
hostname [string]
hosts-allow-include [string]
lcd-display [disabled | enabled]
net-reboot [disabled | enabled]
password-prompt [string]
mgmt-dhcp [dhcpv4 | dhcpv6 | disabled | enabled]
quiet-boot [disabled | enabled]
remote-host [add | delete | replace-all-with] {
[name]... {
options:
addr [IP address]
hostname [string]
}
}
remote-host none
username-prompt [string]

edit global-settings

options:

all-properties
non-default-properties

DISPLAY

list global-settings

list global-settings [option]

show running-config global-settings

show running-config global-settings [option]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the global-settings component to set up the BIG-IP system.

EXAMPLES

```
modify system remote-host add { bigip151 {addr 172.27.226.151 hostname bigip151.saxon.net} }
```

Sets up a remote host named bigip151 with an IP address of 172.27.226.151 and a hostname of bigip151.saxon.net.

```
list global-settings all-properties
```

Displays all of the properties of the global system settings.

OPTIONS

aws-access-key

Amazon Web Services (AWS) supplied access key needed to make secure requests to AWS. The default value is none.

aws-secret-key

Amazon Web Services (AWS) supplied secret key needed to make secure requests to AWS. The default value is none.

aws-api-max-concurrency

Maximum concurrent connections allowed while making Amazon Web Service (AWS) api calls. The default value is 1.

file-blacklist-path-prefix

Specifies the path prefixes that are disallowed for certain commands. The blacklist takes precedence over the whitelist. It is used by the tmsh save/load sys config file command to disallow saving or loading configuration. Example: The path prefix /shared/tmp/ is included both in the whitelist and blacklist. Since, it is present in the blacklist, the configuration cannot be saved or loaded from the /shared/tmp/ location. The paths are specified in braces separated by spaces in quotes. ex: "{/shared/3dns}/shared/bin/".

file-blacklist-read-only-path-prefix

Specifies the read-only path prefixes that are disallowed for certain commands. It is used by the tmsh save/load sys config file command to disallow saving or loading configuration. It is a read-only attribute with value "{/etc/shadow}".

file-whitelist-path-prefix

Specifies the path prefixes that are valid for certain commands. It is used by the tmsh save/load sys config file command for saving or loading configuration. The paths are specified in braces separated by spaces in quotes. ex: "{/var/local/scf}/tmp/ /shared/ /config/".

console-inactivity-timeout

Specifies the number of seconds of inactivity before the system logs off a user that is logged on. The default value is 0 (zero), which means that no timeout is set. The valid range is 0 - 2147483647.

custom-addr

Specifies an IP address for the system. The default value is ::. The host-addr-mode option must be set to custom in order for this setting to take effect.

description

Specifies a user defined description. The default value is no description.

failsafe-action

Specifies the action that the system takes when the switch board fails. The default value is go-offline-restart-tm.

failover-restart-tm

Specifies that when the switch board fails the system restarts the traffic management system and fails over to the other unit in a redundant pair.

go-offline

Specifies that when the switch board fails the system goes offline.

go-offline-restart-tm

Specifies that when the switch board fails the system goes offline and restarts the traffic management system.

reboot

Specifies that after the active cluster fails over to its peer, it reboots while the peer processes the traffic.

restart-all

Specifies that when the switch board fails the system restarts all system services.

file-local-path-prefix

Specifies a list of folder prefixes that can be applied for file objects. This is a space separated list of folder prefixes, contained in curly braces. Example: "{file:///shared/}" or "{file:///fileobjectfolder/} /shared/". By default the folders are "/shared/" and "/tmp/", represented as "{/shared/} {/tmp/}".

gui-audit

Specifies whether or not system GUI log audit messages. If you disable this option, system GUI will not log audit messages. The default value is disabled.

gui-expired-cert-alert

Specifies whether or not system GUI identify in use expired certificates and alert the user. If you

disable this option, system GUI will not monitor in use certificates. The default value is enabled.

gui-security-banner

Specifies whether the system presents on the login screen the text you specify in the `gui-security-banner-text` option. If you disable this option, the system presents an empty frame in the right portion of the login screen. The default value is enabled.

gui-security-banner-text

Specifies the text to present on the login screen when the `gui-security-banner` option is enabled. The default value is Welcome to the BIG-IP Configuration Utility.

Note: To enter a carriage return in the text type Ctrl-V followed by Ctrl-J. Additionally, you must escape special characters, such as a question mark(?), with a back slash.

gui-setup

Enables or disables the Setup utility in the browser-based Configuration utility. The default value is enabled.

Note: When you configure a system using `tmsch`, disable this option. Disabling this option allows the system administrators to use the browser-based Configuration utility without having to run the Setup utility.

host-addr-mode

Specifies the type of host address you want to assign to the system. The default value is management. The options are:

custom

Use this value to specify a custom IP address for the system using the `custom-addr` option.

management

Indicates that the host address is the management port of the system.

state-mirror

Use this value when the host address of the system is shared by the other system in a redundant pair. In case of system failure, the traffic to the other system is routed to this system.

hostname

Specifies a local name for the system. The default value is `bigip1`.

hosts-allow-include

Warning: Do not use this parameter without assistance from the F5 Technical Support team. The system does not validate the commands issued when you use the `hosts-allow-include` option. If you use this option incorrectly, you put the functionality of the system at risk.

lcd-display

Enables or disables the LCD display on the front of the system. The default value is enabled.

net-reboot

Enables or disables the network reboot feature. The default value is disabled.

If you enable this feature and then reboot the system, the system boots from an ISO image on the network, rather than from an internal media drive. Use this option only when you want to install software on the system, for example, for an upgrade or a re-installation.

Note: An enabled value reverts to disabled after you reboot the system a second time.

password-prompt

Specifies the text to present above the password field on the system's login screen.

mgmt-dhcp

Specifies whether the system uses DHCPv4/DHCPv6 clients for acquiring the management interface IP addresses. The option takes 4 possible values: `dhcpv4`, `dhcpv6`, `disabled`, `enabled`. `dhcpv4` and `dhcpv6` options only enable DHCPv4 or DHCPv6 client respectively. `enabled` and `disabled` options enable/disable both DHCPv4 and DHCPv6 clients.

If this option is enabled, manually specified IP addresses for the management interface may be overwritten if the network also contains a DHCP server (for the given IP protocol). If this option is disabled, no DHCP server will be applied to the management interface, however any previously acquired address will still be used. The default value is enabled for VE and disabled for all other platforms. When this option is enabled, manual changes like create/delete on `sys management-ip` will not be allowed. For `dhcpv4/dhcpv6` values, this only applies to the `management-ip` entries matching the IP protocol. For example, for `dhcpv4` value, user can't manually change IPv4 `management-ip` but user can change IPv6 `management-ip`.

quiet-boot

Enables or disables the quiet boot feature. The default value is enabled. When enabled, the system suppresses informational text on the console during the boot cycle.

remote-host

Configures a remote host in the `/etc/hosts` file. The default value is none. You must enter both an IP address and a fully qualified domain name (FQDN) or alias for each host that you want to add to the file.

username-prompt

Specifies the text to present above the user name field on the system's login screen.

SEE ALSO

`edit`, `list`, `modify`, `show`, `tmsch`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2019-02-28 sys global-settings(1)

sys ha-group

NAME

ha-group - Configures the high availability (HA) scoring mechanism for a unit in a traffic group of BIG-IP(r) systems.

MODULE

sys

SYNTAX

Configure the ha-group component within the sys module using the following syntax.

CREATE/MODIFY

```
create ha-group [name]
modify ha-group [name]
options:
  active-bonus [integer]
  app-service [[string] | none]
  clusters none
  clusters [add | delete | modify | replace-all-with] {
    [name] {
      app-service [[string] | none]
      attribute percent-up-members
      threshold [integer]
      minimum-threshold [integer]
      sufficient threshold [integer | all]
      weight [integer]
    }
  }
  description [string]
  [disabled | enabled]
  pools none
  pools [add | delete | modify | replace-all-with] {
    [name] {
      app-service [[string] | none]
      attribute percent-up-members
      threshold [integer]
      minimum-threshold [integer]
      sufficient threshold [integer]
      weight [integer]
    }
  }
  trunks none
  trunks [add | delete | modify | replace-all-with] {
    [name] {
      app-service [[string] | none]
      attribute percent-up-members
      threshold [integer]
      minimum-threshold [integer]
      sufficient threshold [integer]
      weight [integer]
    }
  }
}
```

DISPLAY

```
list ha-group
list ha-group [name]
options:
  all
  all-properties
  current-module
  one-line
```

DELETE

```
delete ha-group [name]
```

DESCRIPTION

You can use the ha-group component to configure a high availability (HA) group that determines the HA scoring mechanism for a unit in a traffic group. This mechanism compares the relative health of the two or more units

in the traffic group and the system with the highest score becomes the active unit. Note Use the attribute ha-group of the traffic group to make the association.

EXAMPLES

```
create ha-group group1 pools add { ftp_pool { attribute percent-up-members weight 70 } }
```

Creates a HA group, named group1, that includes the pool named ftp_pool, and uses the attribute percent-up-members and a weight of 70 to determine the HA score for a unit in a traffic group.

```
list ha-group group1
```

Displays the configuration of the HA group, group1.

OPTIONS

active-bonus

Specifies a number to add to the unit's HA score when the unit is active. This option ensures that the state of a unit is dependent upon the history of its state. The default value is 10 (ten). The range is 0 - 100.

app-service

Specifies the name of the application service to which the object belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

attribute

Specifies an attribute of the component that you want to use for the HA scoring mechanism. Percent-up-members is the only available attribute for HA scoring for the clusters, pools, and trunks options.

clusters

Specifies the clusters that you want to configure for the HA group. You can only configure a cluster on a chassis.

description

User defined description.

[disabled | enabled]

Enables or disables the HA group in the HA table. The default value is enabled.

name Specifies the name of the component that you want to configure. This option is required when you create, modify, or delete a HA group. This option is also required when you configure clusters, pools, or trunks for the HA group.

pools

Specifies the pools that you want to configure for the HA group.

threshold

Deprecated. Use minimum-threshold instead.

minimum-threshold

Specifies the minimum number of up interfaces in a trunk, up pool members in a pool, or up cluster members in a cluster below which the specified component does not contribute to the HA score for the unit. The default value is 0 (zero), which indicates this option is disabled. The value may not exceed the number of members of the trunk, pool, or cluster.

sufficient-threshold

Specifies the sufficient number of up interfaces in a trunk, up pool members in a pool, or up cluster members in a cluster above which the specified component is considered at 100% for its contribution to the HA score for the unit. The default value is all (or 0). The value may not exceed the number of members of the trunk, pool, or cluster.

trunks

Specifies the trunks that you want to configure for the HA group.

weight

The value of this option is multiplied by the percent of up cluster, pool, or trunk members, and is added to the HA score. The default value is 10. The range is 10 - 100.

SEE ALSO

create, delete, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-04-08 sys ha-group(1)

NAME

ha-status - Displays information about the high availability (HA) status of a unit in a redundant pair.

MODULE

sys

SYNTAX

Display information about the ha-status component within the sys module using the following syntax.

DISPLAY

show ha-status

options:

all-properties

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

failure

field-fmt

DESCRIPTION

You can use the ha-status component to display information about the high availability status of a unit in a redundant pair.

EXAMPLES

show ha-status

Display information about the HA status of the unit.

OPTIONS

failure

Display only the objects that present a failure condition.

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2017-05-24 sys ha-status(1)

sys hardware

NAME

hardware - Displays the BIG-IP(r) system hardware.

MODULE

sys

SYNTAX

Display statistics for the hardware component within the sys module using the syntax in the following section.

DISPLAY

show hardware

DESCRIPTION

You can use the hardware component to display information about the hardware.

EXAMPLES

show hardware

Displays hardware information for the system.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

sys host-info

NAME

host-info - Displays statistics about the host.

MODULE

sys

SYNTAX

Configure the host-info component within the sys module using the syntax in the following sections.

DISPLAY

show host-info

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

global

DESCRIPTION

You can use the host-info component to display statistics about the host, including CPU count, active CPU count, processor mode, memory usage, and more.

EXAMPLES

show host-info

Displays host statistics in the system default units.

show host-info raw

Displays raw host statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

sys httpd

NAME

httpd - Configures the HTTP daemon for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the httpd component within the sys module using the following syntax.

CREATE/MODIFY

modify httpd

options:

allow [add | delete | none |replace-all-with] {
hostname or IP address ...

}

auth-name [string]

auth-pam-dashboard-timeout [off | on]

auth-pam-idle-timeout [integer]

auth-pam-validate-ip [off | on]

description [string]

fastcgi-timeout [integer]

hostname-lookup [double | off | on]

include [string]
log-level [alert | crit | debug | emerg | error | info | notice | warn]
redirect-http-to-https [disabled | enabled]
request-header-max-timeout [integer]
request-header-min-rate [integer]
request-header-timeout [integer]
request-body-max-timeout [integer]
request-body-min-rate [integer]
request-body-timeout [integer]
ssl-ca-cert-file [string]
ssl-certchainfile [string]
ssl-certfile [string]
ssl-certkeyfile [string]
ssl-ciphersuite [string]
ssl-include [string]
ssl-protocol [string]
ssl-port [integer]
ssl-verify-client [no | require | optional | optional-no-ca]
ssl-verify-depth [integer]
ssl-ocsp-enable [on | off]
ssl-ocsp-default-responder [string]
ssl-ocsp-override-responder [on | off]
ssl-ocsp-responder-timeout [integer]
ssl-ocsp-response-max-age [integer]
ssl-ocsp-response-time-skew [integer]

edit httpd

options:

all-properties
non-default-properties

DISPLAY

list httpd

list httpd [option name]

show running-config httpd

show running-config httpd [option name]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the httpd component to configure the HTTP daemon for the system.

Important: F5 Networks recommends that users of the Configuration utility exit the utility before changes are made to the system using the httpd component. This is because making changes to the system using this component causes a restart of the httpd daemon. Additionally, restarting the httpd daemon creates the necessity for a restart of the Configuration utility.

EXAMPLES

```
modify httpd { ssl-certfile [string] ssl-certkeyfile [string] }
```

Changes the SSL certificate and the SSL key. Note that when you change the SSL key, you must also change the SSL certificate.

```
modify httpd auth-pam-idle-timeout 43200
```

Sets the PAM idle timeout to half a day (in seconds).

```
modify httpd allow replace-all-with {172.27.0.0/255.255.0.0}
```

Replaces the existing list of hosts that can connect to the httpd daemon with the hosts in the range, 172.27.0.0/255.255.0.0.

OPTIONS

allow

Configures IP addresses and hostnames for the HTTP clients from which the httpd daemon accepts requests. The default value is All.

Warning: Using the value none resets the httpd daemon to allow NO HTTP clients access to the system; therefore, F5 Networks recommends that you do not use the value none.

auth-name

Specifies the name for the authentication realm. The default value is BIG-IP.

auth-pam-dashboard-timeout

Specifies whether idle timeout while viewing the dashboard is enforced or not. The default value is off.

auth-pam-idle-timeout

Specifies the number of seconds of inactivity that can elapse before the GUI session is automatically logged out. The default value is 1200 seconds.

auth-pam-validate-ip

Specifies whether the check for consistent inbound IP for the entire web session is enforced or not. The default value is on.

description

User defined description.

fast-cgitimeout

Specifies, in seconds, the timeout for FastCGI. The default value is 300 seconds.

fips-cipher-version

Read-only field for internal use. Non-zero value indicates that ssl-ciphersuite has been set to FIPS 140-2 compliant defaults. The value 1 indicates that the ciphersuite is "ECDHE-RSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-SHA:ECDHE-RSA-AES256-SHA:ECC" User changes to ssl-ciphersuite will not affect this field. This field is relevant only when FIPS 140-2 compliance is enabled in the license.

hostname-lookup

The default value is off.

include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the include option incorrectly, you put the functionality of the system at risk.

log-level

Specifies the minimum httpd message level to include in the system log. The default value is warn.

redirect-http-to-https

Specifies whether the system should redirect HTTP requests targeted at the configuration utility to HTTPS. The default value is disabled.

request-header-max-timeout

Specifies, in seconds, the maximum time allowed to receive all of the request headers, if the request-header-min-rate option is used, in which case the timeout is extended as more data arrives. Ignored if request-header-min-rate is not used. A value of 0 means no limit. The default value is 40.

request-header-min-rate

Specifies, in bytes per second, the minimum average rate at which the request headers must be received. A value of 0 means no limit. The default value is 500.

request-header-timeout

Specifies, in seconds, the time allowed to receive all of the request headers. A value of 0 means no limit. If you use the request-header-min-rate option, this represents the initial value for the timeout, which will be extended as more data arrives. The default value is 20.

Warning: This includes the time needed to complete the initial SSL handshake. If the user's browser is configured to query certificate revocation lists and the CRL server is not reachable, the initial SSL handshake may take a significant time until the browser gives up waiting for the CRL.

request-body-max-timeout

Specifies, in seconds, the maximum time allowed to receive all of the request body, if the request-body-min-rate option is used, in which case the timeout is extended as more data arrives. Ignored if request-body-min-rate is not used. A value of 0 means no limit. The default value is 0.

request-body-min-rate

Specifies, in bytes per second, the minimum average rate at which the request body must be received. A value of 0 means no limit. The default value is 500.

request-body-timeout

Specifies, in seconds, the time allowed for reading all of the request body. This includes the time needed to do any SSL renegotiation. A value of 0 means no limit. If you use the request-body-min-rate option, this represents the initial value for the timeout, which will be extended as more data arrives. The default value is 60.

ssl-ca-cert-file

Specifies the name of the file that contains the SSL Certificate Authority (CA) certificate file. The default value is none.

ssl-certchainfile

Specifies the name of the file that contains the SSL certificate chain. The default value is none.

ssl-certfile

Specifies the name of the file that contains the SSL certificate. The default value is /etc/httpd/conf/ssl.crt/server.crt.

Note that the path to the file must start with either /etc/httpd/conf/ssl.crt/ or /config/httpd/conf/ssl.crt/, unless the path is a relative path. If the path is a relative path, then it must start with conf/ssl.crt/.

ssl-certkeyfile

Specifies the name of the file that contains the SSL certificate key. The default value is /etc/httpd/conf/ssl.key/server.key.

Note that the path to the file must start with either /etc/httpd/conf/ssl.key/ or /config/httpd/conf/ssl.key/, unless the path is a relative path. If the path is a relative path, then it must start with conf/ssl.key/.

When you change the key file, you must also change the certificate file. For example, use the following command sequence to change the key: `modify httpd { ssl-certfile [string] ssl-certkeyfile [string] }`

ssl-ciphersuite

Specifies the ciphers that the system uses. The default value is "ECDHE-RSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-SHA:ECDHE-RSA-AES256-SHA:EC

ssl-include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the `ssl-include` option incorrectly, you put the functionality of the system at risk.

ssl-protocol

The list of SSL protocols to accept on the management console. A space-separated list of tokens in the format accepted by the Apache `mod_ssl` `SSLProtocol` directive.

The default value is all `-SSLv2 -SSLv3`.

ssl-port

The SSL port to run the management console. It is a number in the range of 1 and 65535.

The default value is 443.

ssl-ocsp-default-responder

Specifies the default responder URI for OCSP validation. The default is `http://localhost.localdomain`. The value for the default responder should always be preceded with `http://`.

ssl-ocsp-enable

Specifies OCSP validation of the client certificate chain. The default is off.

ssl-ocsp-override-responder

Specifies the force use of default responder URI for OCSP validation. The default is off.

ssl-ocsp-responder-timeout

Specifies the maximum allowable time in seconds for OCSP response. The default is 300 seconds.

ssl-ocsp-response-max-age

Specifies the maximum allowable age ("freshness") for OCSP responses. The default value (-1) does not enforce a maximum age, which means that OCSP responses are considered valid as long as their `nextUpdate` field is in the future.

ssl-ocsp-response-time-skew

Specifies the maximum allowable time skew in seconds for OCSP response validation. The default is 300 seconds.

ssl-verify-client

Specifies if the client certificate needs to be verified for SSL session establishment. The default is no.

ssl-verify-depth

Specifies maximum depth of CA certificates in client certificate verification. The default is 10.

SEE ALSO

`edit`, `list`, `modify`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2018-10-30 sys httpd(1)

sys hypervisor-info

NAME

`hypervisor-info` - Used inside of a vCMP guest to display proposed configuration information passed in from the vCMP hypervisor.

MODULE

`sys`

SYNTAX

Access the `hypervisor-info` component within the `sys` module using the syntax in the following sections.

DISPLAY

`show hypervisor-info`
options:

field-fmt

DESCRIPTION

You can use the `hypervisor-info` component to display vCMP guest configuration information proposed by the vCMP hypervisor. Note that this component will only display information when used from inside a vCMP guest. On any other BIG-IP system, the `show` command will produce no output.

These values will override the default values for any of the corresponding configuration items inside the guest:

Proposed Address - The management IP, as configured via `tmsh sys management-ip` on appliances and via `tmsh sys cluster` on clusters.

Proposed Gateway - The default gateway, as configured via `tmsh sys management-route`.

Proposed Hostname - The hostname, as configured via `tmsh sys global-settings`.

Proposed Netmask - The management netmask. See "Proposed Address" above.

EXAMPLES

```
show hypervisor-info
```

Displays hypervisor configuration information in default units.

OPTIONS

For information about the options that you can use with the command `show`, see `help show`.

SEE ALSO

`show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013-2014. All rights reserved.

BIG-IP 2016-03-14 sys hypervisor-info(1)

sys icall event

NAME

`event` - Generate an Event on the BIG-IP(r) system.

MODULE

`sys icall`

SYNTAX

Generate the event component within the `sys icall` module using the syntax shown in the following sections.

GENERATE

```
generate event
options:
name [string]
context {
{
name [string]
value [string]
}
}
```

DESCRIPTION

You may use the `generate event` command to construct a free-form Event in the system which will be sent to interested Event Handlers.

EXAMPLES

```
generate event name EMPLOYEE context { { first_name Sam } { last_name Shepard } }
```

Construct an event named "EMPLOYEE" that contains two pieces of information as name/value pairs. An Event Handler must be subscribed to the event by the name "EMPLOYEE" or by both event name and all the contexts in a context group.

OPTIONS

`context`

Specifies a set of name/value pairs that convey the information of the Event.

`name` The Events name; does not have to be unique, but may not be empty.

SEE ALSO

create, delete, edit, list, modify, show, sys ical event-handler, sys ical script, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys ical event(1)

sys ical handler periodic

NAME

periodic - Make or configure a periodic handler for the BIG-IP(r) system.

MODULE

sys ical handler

SYNTAX

Modify the periodic component within the sys ical handler module using the syntax shown in the following sections.

CREATE/MODIFY

create periodic [name]

modify periodic [name]

options:

arguments {

{

name [string]

value [string]

}

...

}

description [string]

first-occurrence [date/time]

interval [integer]

last-occurrence [date/time]

script [script name]

status [active | inactive]

mv periodic [[[source-name] [destination-name]] |

[[name] to-folder [folder-name]] |

[[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list periodic

list periodic [[[name] | [glob] | [regex]] ...]

show periodic

show periodic [[[name] | [glob] | [regex]] ...]

DELETE

delete periodic [name]

DESCRIPTION

You can create a periodic handler to run scripts automatically based on clock time.

EXAMPLES

```
create periodic my_handler1 script script1 first-occurrence now+1h interval 45 arguments { { name user value j.han } { name role value manager } }
```

Create a new periodic handler that will execute script1 every 45 seconds. The handler will wait one hour before beginning, but continue to execute indefinitely. Each 45 seconds, when the script executes, the provided arguments will be passed into the script as `EVENT::context() data`.

```
mv periodic /Common/my_periodic to-folder /Common/my_folder
```

Moves a periodic ical handler named my_periodic to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

arguments

Specifies a set of name/value pairs that will be passed to the script at the start of each execution on each interval.

The use of arguments is optional and may be changed at any time.

description
A user defined description of the item.

first-occurrence
A specific date and time for this handler to begin executing. If not specified, the current date and time of creation will be used.

interval
The number of seconds between each time this handler should execute.

last-occurrence
A specific date and time for this handler to stop executing. If not specified, the script will run indefinitely.

script
The iCall Tcl script the handler when execute at each time interval. Note that this script must be an object in sys icall script; a cli script will not work.

status
Specify either active or inactive. Active is the default value.

When the handler status is active, the handler accepts events and executes the script as expected. However, when the status is inactive, the handler will no longer accept incoming events and the script will not execute. Use the inactive status when you wish to keep the handler as a configuration item and do not wish to delete it, but also do not wish the handler to run.

to-folder
A periodic icall handler can be moved to any folder under /Common, but configuration dependencies may restrict it from moving out of /Common.

SEE ALSO
create, delete, edit, list, modify, mv, show, sys icall event, sys icall script, tmsb

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys icall handler periodic(1)

sys icall handler perpetual

NAME
perpetual - Make or configure a perpetual handler for the BIG-IP(r) system.

MODULE
sys icall handler

SYNTAX
Modify the perpetual component within the sys icall handler module using the syntax shown in the following sections.

CREATE/MODIFY
create perpetual [name]
modify perpetual [name]
options:
description [string]
script [script name]
status [active | inactive | suspend]
subscriptions [add | delete | modify | replace-all-with] {
[subscription name] {
options:
event-name [event name]
filters [add | delete | modify | replace-all-with] {
[filter name] {
options:
value [string]
match-algorithm [accept-all | exact | glob | regex | subnet]
}
}
}
}
restart perpetual [name]
start perpetual [name]
stop perpetual [name]

mv perpetual [[source-name] [destination-name]] |

```
[ [name] to-folder [folder-name] ] |
[ [name...name] to-folder [folder-name] ] ]
```

options:
to-folder

DISPLAY

```
list perpetual
list perpetual [ [ [name] | [glob] | [regex] ] ... ]
show perpetual
show perpetual [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete perpetual [name]
```

DESCRIPTION

You can create a perpetual handler to run continuously executing code and to receive events by specifying subscriptions.

EXAMPLES

```
create perpetual my_handler1 script script1 subscriptions add { sub1 { event-name LTM_POOL_UP } }
```

Creates a new perpetual handler run the program defined in "script1". Anytime an event called "LTM_POOL_UP" is generated in the system, a copy will be sent to my_handler1.

```
mv perpetual /Common/my_perpetual to-folder /Common/my_folder
```

Moves a perpetual icall handler named my_perpetual to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

description

A user defined description of the item.

script

The iCall Tcl script the handler will execute upon creation. The user is responsible for creating a script with perpetual execution. If the script is changed, the handler will not change its executing code until the handler is restarted or put into inactive and then active status.

Note that this script must be an object in sys icall script; a cli script will not work.

status

Specify active, inactive, or suspend. Active is the default value.

Inactive status indicated that the handler is to no longer execute and to no longer receive events. The handler's state is lost and all pending events are deleted. Use this status to eliminate a handler in the system but to keep its information stored.

The handler may also be set to suspend which will keep the handler script executing, but the system will send no new events to the handler. Events waiting to be processed remain in queue.

subscriptions

Create one or more subscription items to specify the conditions of this handler's execution. The handler subscribes generally to events by the event name, and specifically to data by using filters. The use of filters is optional.

The handler will be sent events by the system as defined by the subscription property, but the code inside the handler must use EVENT::get_next function in order to receive the data into the handler. See sys icall script for more information.

to-folder

A perpetual icall handler can be moved to any folder under /Common, but configuration dependencies may restrict it from moving out of /Common.

SEE ALSO

create, delete, edit, list, modify, mv, show, sys icall event, sys icall script, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys icall handler perpetual(1)

sys icall handler triggered

NAME

triggered - Make or configure an event-triggered handler for the BIG-IP(r) system.

MODULE

sys icall handler

SYNTAX

Modify the triggered component within the sys icall handler module using the syntax shown in the following sections.

CREATE/MODIFY

```
create triggered [name]
modify triggered [name]
options:
  description [string]
  script [script name]
  status [active | inactive]
  subscriptions [add | delete | modify | replace-all-with] {
    [subscription name] {
      options:
        event-name [event name]
        filters [add | delete | modify | replace-all-with] {
          [filter name] {
            options:
              value [string]
            match-algorithm [accept-all | exact | glob | regex | subnet]
          }
        }
      }
    }
  }
}
```

```
mv triggered [ [ [source-name] [destination-name] ] |
  [ [name] to-folder [folder-name] ] |
  [ [name...name] to-folder [folder-name] ] ]
options:
  to-folder
```

DISPLAY

```
list triggered
list triggered [ [ [name] | [glob] | [regex] ] ... ]
show triggered
show triggered [ [ [name] | [glob] | [regex] ] ... ]
```

DELETE

```
delete triggered [name]
```

DESCRIPTION

You can create a triggered handler to automatically run a script when a specified event occurs.

EXAMPLES

```
create triggered my_handler1 script script1 subscriptions add { pools { event-name LTM_POOL_UP filters add {
pool_name { value pool1 } node_name { value node1 } } } }
```

Creates a new triggered handler that will execute the script called "script1" when an event called "LTM_POOL_UP" is generated in the system and contains the contexts { pool_name, pool1 } and { node_name, node1 }.

```
mv triggered /Common/my_triggered to-folder /Common/my_folder
```

Moves a triggered icall handler named my_triggered to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

description

A user defined description of the item.

script

The iCall Tcl script the handler will execute when invoked by an appropriate event. Note that this script must be an object in sys icall script; a cli script will not work.

status

Specify either active or inactive. Active is the default value.

When the handler status is active, the handler accepts events and executes the script as expected. However, when the status is inactive, the handler will no longer accept incoming events and the script will not execute. Use the inactive status when you wish to keep the handler as a configuration item and do not wish to delete it, but also do not wish the handler to run.

subscriptions

Specify one or more subscriptions to define the conditions of this handler's execution. The handler subscribes generally to events by the event name, and specifically to data by using filters. The use of filters is optional.

A handler that specifies more than one subscription will execute when any one subscription is matched to an event.

to-folder

A triggered icall handler can be moved to any folder under /Common, but configuration dependencies may restrict it from moving out of /Common.

SEE ALSO

create, delete, edit, list, modify, mv, show, sys icall event, sys icall script, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys icall handler triggered(1)

sys icall istats-trigger

NAME

istats-trigger - Configure an iStats trigger to generate a user defined event for the iCall feature on the BIG-IP(r) system.

MODULE

sys icall

SYNTAX

Modify the istats-trigger component within the sys icall module using the syntax shown in the following sections.

CREATE/MODIFY

create istats-trigger [name]

modify istats-trigger [name]

options:

description [string]

duration [integer]

event-name [string]

istats-key [string]

range-max [integer]

range-min [integer]

repeat [integer]

mv istats-trigger [[[source-name] [destination-name]] |

[[name] to-folder [folder-name]] |

[[name...name] to-folder [folder-name]]]

options:

to-folder

DISPLAY

list istats-trigger

list istats-trigger [[[name] | [glob] | [regex]] ...]

DELETE

delete istats-trigger [name]

DESCRIPTION

You can create an istats-trigger to automatically generate a Control Plane iRules event under the conditions specified in the properties.

EXAMPLES

mv istats-trigger /Common/my_istats_trigger to-folder /Common/my_folder

Moves an istats-trigger named my_istats_trigger to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

description

A user defined description of the item.

duration

Duration in seconds. The value "0" means trigger instantly when in range.

event-name

The name of the event that will be generated.

istats-key

Specify the items and thresholds to define when this istats-trigger will generate an event.

range-max

Trigger event only if value is less-than-or-equal to range-max.

range-min

Trigger event only if value is greater-than-or-equal to range-min. Note that if 0 is included in the

specified range, then the iStats key must be explicitly initialized with "istats set [key] 0" in order for the trigger to fire.

repeat
Repeat interval in seconds. The value "none" means do not resend the event unless the value falls outside the range and then re-enters it.

to-folder
An istats-trigger can be moved to any folder under /Common, but configuration dependencies may restrict it from moving out of /Common.

SEE ALSO

create, delete, edit, list, modify, mv, show, sys icall event, sys icall event-handler, sys icall script, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys icall istats-trigger(1)

sys icall publisher

NAME

publisher - Show the services publishing events on a BIG-IP(r) system

MODULE

sys icall

SYNTAX

Show the available publishers within the sys icall module using the syntax shown in the following sections.

DISPLAY

show publisher [field-fmt]
show publisher [[[name] | [glob] | [regex]] ...] [field-fmt]

DESCRIPTION

This command lets you display the publishers on the system, as well as the events that they publish and the contexts that those events are guaranteed to contain.

By default these are shown in a tabular form; use the field-fmt option to show them in a format similar to listing other objects in tmsh.

If a published event includes no contexts, then a single line will be shown with a - in the context column. If a publisher publishes no events, then a single line will be shown with a - in the event column.

OPTIONS

field-fmt

By default, the events will be shown in a tabular format. This overrides the command to print the publishers in object format like the list command does for other objects.

SEE ALSO

show, tmsh, sys icall event, sys icall handler periodic, sys icall handler perpetual, sys icall handler triggered, sys icall istats-trigger, sys icall script

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2014-06-27 sys icall publisher(1)

sys icall script

NAME

script - Manage a Tcl script used by handlers during execution on the BIG-IP(r) system.

MODULE

sys icall

SYNTAX

Manage the script component within the sys icall module using the syntax shown in the following sections.

CREATE/MODIFY/EDIT

```
create script [name]
modify script [name]
edit script [name]
options:
  definition
  description [string]
  events [add | delete | modify | replace-all-with] {
    [event name] {
      contexts [add | delete | modify | replace-all-with] {
        [context name]
      }
    }
  }
}
```

```
mv script [ [source-name] [destination-name] ] |
  [ [name] to-folder [folder-name] ] |
  [ [name...name] to-folder [folder-name] ] ]
options:
  to-folder
```

DISPLAY

```
list script
list script [name]
```

DELETE

```
delete script [name]
```

Note: You must remove all references to the icall script before deletion.

DESCRIPTION

You can use this script component to manage Tcl scripts which are used by event handlers upon execution.

Caution: If you add a handler to a shared configuration on a set of BIG-IP appliances, then care must be used in making changes to configuration items. A handler's script which makes config changes on more than one device may cause inconsistencies that must be manually resolved.

EXAMPLES

```
create script my_script1
```

Create a new icall script item called "my_script1". Upon pressing enter, the user will enter the text editor in order to edit the Tcl script. Note that this configuration item may only be modified while in the edit view.

```
mv script /Common/my_script to-folder /Common/my_folder
```

Moves an icall script named my_script to a folder named my_folder, where my_folder has already been created and exists within /Common.

OPTIONS

definition
Holds the Tcl code.

description
User defined description.

events
Register events with the system that this script creates.

to-folder
An icall script can be moved to any folder under /Common, but configuration dependencies may restrict it from moving out of /Common.

EVENT ACCESSORS

In addition to all the tmsh:: commands provided by the system to use in the Tcl scripts (please see help cli script), the commands below are provided to access event specific information.

Hint: When you use a tmsh:: command, call it inside of Tcl catch to receive any error messages returned, and to allow the script to exit gracefully if needed. Without Tcl catch, the script may crash and end the process.

The following Tcl variables may be used in triggered handlers. (The \$ is not part of the variable name but is the lookup operator for the Tcl variable.):

\$EVENT::context([name])

An array variable containing the value of each context, keyed by the context name.

\$EVENT::creation_time

The date and time the event was generated.

\$EVENT::event_name

The name of the event that was generated.

`$EVENT::handler_name`

The name of the event handler that matched the event being handled.

`$EVENT::script_name`

The name of the currently running script.

For use in perpetual handlers:

`EVENT::get_next [-timeout [milliseconds]]`

The timeout parameter is optional. If the timeout is set, then `EVENT::get_next` will return 0 if no event matches before the timeout hits. Otherwise, the `EVENT::get_next` will return 1, and the above variables in the `EVENT::` namespace will be replaced with the data from the new event.

SCRIPT EXAMPLES

The following script will print out all the information of an event.

```
puts "*** start of event ***"

foreach var [info vars EVENT::*] {
  set varname [namespace tail $var]
  if { [array exists $var] } {
    puts "$varname: "
    foreach { k v } [array get $var] { ;#k = key v = value
      puts "$k:$v"
    }
  } else {
    puts "$varname: [set $var]"
  }
}
```

The next script will allow events to hold bash commands and have the script execute them. The script would be required to run inside an event handler that subscribed to the appropriate event and filtered on the words "utility" and "arguments".

```
set bash_cmd $EVENT::context(utility) append bash_cmd " " $EVENT::context(arguments)
```

```
if { [catch { exec /bin/bash -c $bash_cmd } result] } {
  puts "error executing bash command: $bash_cmd" } else {
  puts $result }
```

SEE ALSO

`cli script`, `create`, `delete`, `edit`, `list`, `modify`, `show`, `sys icall event`, `sys icall event-handler`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys icall script(1)

sys icmp-stat

NAME

`icmp-stat` - Displays and resets ICMP statistics on the BIG-IP system.

MODULE

sys

SYNTAX

Configure the `icmp-stat` component within the `sys` module using the syntax in the following section.

MODIFY

```
reset-stats icmp-stat
```

DISPLAY

```
show icmp-stat
```

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the `icmp-stat` component to display and reset ICMP statistics. The statistics you can view are standard ICMP statistics, including ICMPv4 packets and errors, and ICMPv6 packets and errors.

OPTIONS

For information about the options that you can use with the command `show`, see `help show`.

For information about the options that you can use with the command `reset-stats`, see `help reset-stats`.

SEE ALSO

`reset-stats`, `show`, `sys icmp-stat`, `tms`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 sys icmp-stat(1)

sys iconcontrol-soap

NAME

`iconcontrol-soap` - Configures the iControl SOAP daemon for the BIG-IP(r) system.

MODULE

`sys`

SYNTAX

Configure the `iconcontrol-soap` component within the `sys` module using the following syntax.

CREATE/MODIFY

`modify iconcontrol-soap`

options:

```
allow [add | delete | none |replace-all-with] {  
  All or IP address ...  
}
```

`edit iconcontrol-soap`

options:

```
all-properties  
non-default-properties
```

DISPLAY

`list iconcontrol-soap`

`list iconcontrol-soap [option name]`

`show running-config iconcontrol-soap`

`show running-config iconcontrol-soap [option name]`

options:

```
all-properties  
non-default-properties  
one-line
```

DESCRIPTION

You can use the `iconcontrol-soap` component to configure the iControl SOAP for the system.

Important: F5 Networks recommends that users of the Configuration utility exit the utility before changes are made to the system using the `iconcontrol-soap` component. This is because making changes to the system using this component causes a restart of the `iconcontrol-soap` daemon. Additionally, restarting the `iconcontrol-soap` daemon creates the necessity for a restart of the Configuration utility.

EXAMPLES

```
modify iconcontrol-soap allow replace-all-with {9.9.9.9}
```

Reduces the allowed IP address that can access iControl SOAP to 9.9.9.9

OPTIONS

`allow`

Configures IP addresses for iControl SOAP clients from which the `iconcontrol-soap` daemon accepts requests. The value may be either a full IP address or a Perl Compatible Regular Expression to allow connections from a specific subnet. The default value is `All`.

Warning: Using the value `none` resets the `iconcontrol-soap` daemon to allow all iControl SOAP clients access to the system; therefore, F5 Networks recommends that you do not use the value `none`.

SEE ALSO

`edit`, `list`, `modify`, `show`, `tms`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose

other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2015. All rights reserved.

BIG-IP 2015-07-14 sys iconcontrol-soap(1)

sys integrity status-check

NAME
status-check - Checks and verifies the integrity of the BIG-IP BIOS and kernel using the tpm-status utility.

MODULE
sys integrity

SYNTAX
Run a check on the system's integrity status in the sys integrity module by using the syntax below.

RUN
run status-check
options:
-h Print help for the tpm-status command.
-c Attempt to continue after reaching an error.
-o Output the PCR measurements to a given file after measurement. The file will be overwritten if pre-existing.
-a Append the PCR measurements to a given file after measurement.
-v [0|1|2|3] Modify the verbosity of the tpm-status utility from 0 (not verbose) to 3 (very very verbose). Default value: 0.
-q Generates a file containing system information needed for Remote Attestation. File is located at /var/log/pcrs.json.
Note: Options -o and -a are mutually exclusive. Files may only be written to areas allowed by tmsh.

DESCRIPTION
You can use the status-check command to check the system's integrity status.

EXAMPLES
run status-check

Runs the tpm-status utility and returns the integrity status. It shall be one of VALID, INVALID, UNAVAILABLE, or PENDING.

run status-check -v 3

Runs the tpm-status utility with very very verbose output then returns the integrity status.

SEE ALSO
run, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2017. All rights reserved.

BIG-IP 2018-05-07 sys integrity status-check(1)

sys internal-proxy

NAME
internal-proxy - Configuration of the internal proxy.

MODULE
sys

SYNTAX
Configure the internal-proxy component within the sys.internal-proxy module using the syntax shown in the following sections. This object is used by other services such as CRL cert-validator, for defining how to send/receive the outbound traffic for the service.

CREATE/MODIFY
create internal-proxy [name]
modify internal-proxy [name]
options:
description [string]
dns-resolver [name]

proxy-server-pool [name]
route-domain [name]

DISPLAY
list internal-proxy [name]

DELETE
delete [all | [name]]
options:
recursive

DESCRIPTION
You can use the internal-proxy component to configure a custom internal proxy.

EXAMPLES
create sys internal-proxy my_intp dns-resolver my_dnsr

Creates an internal proxy named my_intp using the DNS resolver my_dnsr.

OPTIONS
dns-resolver
Specifies the DNS resolver object used for sending out the traffic.

proxy-server-pool
Specifies the proxy server pool used for reaching servers.

route-domain
Specifies the route domain for for reaching servers using HTTP forward proxy.

SEE ALSO
sys crypto cert-validator cri

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2017. All rights reserved.

BIG-IP 2017-12-11 sys internal-proxy(1)

sys ip-address

NAME
ip-address - Displays the IP addresses currently associated with a configuration object on a BIG-IP(r) system.

MODULE
sys

SYNTAX
Display the IP addresses associated with a BIG-IP system configuration object using the syntax in the following section.

DISPLAY
show ip-address
options:
[all-properties | field-fmt]

DESCRIPTION
You can use the ip-address component to display the location on the BIG-IP system of the IP addresses associated with a configuration object. The system displays the following information:

Entry
Displays the IP address and any associated configuration. For example, for a Local Traffic Manager pool member, the entry is the member's IP address and port number, 10.1.1.1:80.

Component
Displays the type of component associated with the IP address. For example, for a Local Traffic Manager pool, the entry is ltm pool.

Object-ID
Displays the name of a configuration object associated with the IP address. For example, for a Local Traffic Manager pool named my_pool, the entry is my_pool.

Property
When you specify the all-properties option, displays the name of the property that contains the IP address value. Note that if the IP address is an object identifier the system displays n/a.

EXAMPLES

show ip-address

Displays the IP addresses currently associated with a BIG-IP system configuration object.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-11 sys ip-address(1)

sys ip-stat

NAME

ip-stat - Displays and resets IP statistics on the BIG-IP system.

MODULE

sys

SYNTAX

Configure the ip-stat component within the sys module using the syntax in the following section.

MODIFY

reset-stats ip-stat

DISPLAY

show ip-stat

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the ip-stat component to display and reset IP statistics. The statistics you can view are standard IP statistics, including IPv4 and IPv6 packets, fragments, fragments reassembled, and errors.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, sys ip-stat, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 sys ip-stat(1)

sys ipfix destination

NAME

destination - Displays or resets statistics for IPFIX log destinations.

MODULE

sys ipfix

SYNTAX

Specify the destination within the sys ipfix module using the syntax in the following section.

MODIFY

reset-stats destination [name]

DISPLAY
show destination [name]

DESCRIPTION

You can use the destination component to display IPFIX destination statistics, like Templates and Data Record counts. You can also reset the IPFIX destination statistics to zero at any time.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, sys log-config destination ipfix, sys ipfix irules, sys ipfix element, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-05-27 sys ipfix destination(1)

sys ipfix element

NAME

element - Configures element for IPFIX logging.

MODULE

sys ipfix

SYNTAX

Configure the IPFIX component within the sys ipfix module using the syntax shown in the following sections.

CREATE/MODIFY

create element [name]

modify element [name]

options:

all

app-service [[string] | none]

data-type

[boolean | datetime-microseconds |
datetime-milliseconds | datetime-nanoseconds |
datetime-seconds | float32 | float64 |
ipv4-address | ipv6-address | macaddress |
octetarray | signed16 | signed32 | signed64 |
signed8 | string |
unsigned16 | unsigned32 |
unsigned64 | unsigned8]

description [string]

enterprise-id [integer]

id [integer]

size [integer]

DISPLAY

list element

list element [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete element [name]

Note: Default elements loaded at boot time may not be deleted.

DESCRIPTION

You can use the element component to configure elements for the IPFIX logging interface.

EXAMPLES

create element myelement id 345 enterprise-id 543 data-type string size 128

Creates a element named myelement with element id 345, enterprise-id 543, data-type string and data size of 128.

delete element myelement

Deletes the element named myelement.

list element myelement

Displays properties of the element named myelement.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the element belongs. The default value is none. **Note:** If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the element. Only the application service can modify or delete the element.

data-type

Specifies the data type of the element.

description

User defined description.

enterprise-id

Specifies the enterprise-id for the IPFIX element being configured between 0 and 4294967295. An enterprise id value of 0 is used to define standardized IANA Information Elements.

id Specifies the element id for the IPFIX element being configured between 1 and 65535. Values greater than 32767 will be considered NETFLOW-only Information Elements; since the high bit of the 16-bit value is set for those values.

size Specifies the IPFIX element data size between 1 and 1900 for data-types octetarray and string. The default is 0; and means variable for these two data-types.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique alphanumeric name for the component. Preferably camel casing. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-05-28 sys ipfix element(1)

sys ipfix irules

NAME

irules - Displays or resets statistics for irules that use IPFIX logging destinations.

MODULE

sys ipfix

SYNTAX

Specify the irules component within the sys ipfix module using the syntax in the following section.

MODIFY

reset-stats irules

DISPLAY

show irules

DESCRIPTION

You can use the irules component to display a global set of statistics for the iRules that use IPFIX destinations, like memory allocation and outstanding counts for templates, messages and destinations. You can also reset the IPFIX iRules statistics to zero at any time.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command `reset-stats`, see `help reset-stats`.

SEE ALSO

`reset-stats`, `show`, `sys log-config destination ipfix`, `sys ipfix destination`, `sys ipfix element`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2014-05-27 sys ipfix irules(1)

sys iprep-status

NAME

`iprep-status` - Displays the status of an IP reputation database. In the BIG-IP(R) Configuration Utility, this database is referred to as the IP Address Intelligence database.

MODULE

sys

SYNTAX

Display information about the `iprep-status` component within the `sys` module using the following syntax.

DISPLAY

```
show iprep-status
options:
  current-module
  field-fmt
  running-config
```

DESCRIPTION

You can use the `iprep-status` component to display status information about the IP reputation database. The reputation database (referred to as IP Address Intelligence in the Config Utility) is available from third-party vendors. An IP intelligence database is a list of IP addresses that have a questionable reputation. The status information returned includes:

- the date and time that the BIG-IP system last contacted the vendor server
- the date and time that the BIG-IP system last received an update
- the total number of IP address in the database
- the number of IP addresses in the most recent update

Note: When the system has an IP Intelligence license and the database variable `db iprep.autoupdate` is enabled (default), the database is automatically downloaded and stored in the binary file:

```
/var/lpRep/F5IpRep.dat
```

The database contains information that maps IP addresses or ranges of IP addresses to one or more reputation categories. After every update, the `lpRep` data file is loaded from disk into the running configuration.

EXAMPLES

```
show iprep-status
```

Displays current status information for the IP reputation database.

OPTIONS

For information about the options that you can use with the command `show`, see `help show`.

SEE ALSO

`show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys iprep-status(1)

sys license

NAME

license - Manage the BIG-IP(r) system licensing information.

MODULE

sys

SYNTAX

Manage the license component within the sys module using the syntax in the following section.

INSTALL

install license

options:

add-on-keys { [key] ... }

license-server [[host name] | [IP address]]

license-server-port [number]

no-certificate-update

registration-key [key]

show-difference

verbose

DISPLAY

show license

options:

detail

REVOKE

revoke license

options:

license-server [[host name] | [IP address]]

license-server-port [number]

verbose

DESCRIPTION

You can use the license component to do the following:

Display detailed licensing and version information for the system, including the registration key, licensing dates, platform ID, suggested service check date, and the installed active modules.

Install and update the system license.

Revoke the existing license (only for VE instances; the registration key may subsequently be re-used to relicense this or another VE instance).

EXAMPLES

show license

Displays the system software licensing information.

show license detail

Displays the system software licensing information, including optional modules and active features.

install license

Reactivate an existing license.

revoke license

System returns to an unlicensed state; the previous registration key may be re-used on this or another VE.

OPTIONS

add-on-keys

Specifies additional feature modules to be included in the license. If add-on keys are not specified the system will use the add-on keys in the current license file.

license-server

Specifies the host name or IP address of the license server. The default value is 65.61.115.202 (activate.f5.com).

license-server-port

Specifies the IP port of the license server. The default value is 443.

no-certificate-update

Do not perform the certificate update check when contacting the Licensing Server.

registration-key

Specifies the license registration key. If the registration key is not specified the system will use the registration key in the current license file.

show-difference

Displays a comparison between the existing license and the pending license, and prompts to allow the user to keep the existing license or install the pending license.

verbose

Display status as the license is being installed.

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2017-04-05 sys license(1)

sys log-config destination alertd

NAME

alertd - Modify the AlertD destination.

MODULE

sys log-config destination

SYNTAX

Modify the AlertD component within the sys log-config destination module using the syntax shown in the following sections.

MODIFY

modify alertd [name]

options:

all

description [string]

DISPLAY

list alertd

list alertd [[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

Note: There is only one AlertD destination, alertd; this destination cannot be created or deleted.

DESCRIPTION

You can use this destination component to modify the AlertD destination for the common logging interface. There is only one AlertD destination; it cannot be deleted. You can use this destination to send logs directly to the AlertD daemon, bypassing syslog-ng.

EXAMPLES

list alertd alertd

Displays properties of the AlertD destination.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the modify command.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2014. All rights reserved.

sys log-config destination arcsight

NAME

arcsight - Formats incoming logs into the ArcSight format for delivery by a forwarding destination.

MODULE

sys log-config destination

SYNTAX

Configure the ArcSight component within the sys log-config destination module using the syntax shown in the following sections.

CREATE/MODIFY

```
create arcsight [name]
modify arcsight [name]
options:
  all
  app-service [[string] | none]
  description [string]
  forward-to [string]
```

DISPLAY

```
list arcsight
list arcsight [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

DELETE

```
delete arcsight [name]
```

Note: You must remove all references to a destination before you can delete the destination. Default destinations may not be deleted.

DESCRIPTION

You can use this destination component to create ArcSight formatting destinations for the common logging interface. ArcSight log destinations currently only deliver log messages from the Network Firewall Module or the Application Security Module.

EXAMPLES

```
create arcsight my_dest forward-to another_dest
```

Creates an ArcSight destination named `my_dest` which forwards to another destination `another_dest`. `another_dest` must be a Local Syslog, Local Database, Remote Syslog, or Remote High Speed Log destination.

```
delete arcsight my_dest
```

Deletes the destination named `my_dest`. Destinations cannot be deleted when in use by a publisher.

```
list arcsight my_dest
```

Displays properties of the destination named `my_dest`.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the destination belongs. The default value is `none`. **Note:** If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the destination. Only the application service can modify or delete the destination.

description

User defined description.

forward-to

Specifies a Local Syslog, Local Database, Remote Syslog, or Remote High Speed Log destination. This is required for the create and modify commands.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 sys log-config destination arcsight(1)

sys log-config destination ipfix

NAME

ipfix - Formats log messages into IPFIX messages and sends them to a specified pool of IPFIX Collectors

MODULE

sys log-config destination

SYNTAX

CREATE/MODIFY

create ipfix [name]

modify ipfix [name]

options:

all

app-service [[string] | none]

description [string]

pool-name [string]

protocol-version [ipfix | netflow-9]

template-delete-delay [integer]

template-retransmit-interval [integer]

transport-profile [profile name]

serverssl-profile [profile name]

DISPLAY

list ipfix

list ipfix [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete ipfix [name]

Note: You must remove all references to a destination before you can delete the destination. Default destinations may not be deleted.

DESCRIPTION

You can use this destination component to create IPFIX forwarding destinations for the common logging interface.

The IPFIX protocol is designed for logging IP-transmission events. RFC 5101

() specifies the protocol, and RFC 5102

() describes the information model for IPFIX logs. IPFIX logs are raw,

binary-encoded strings with their fields and field lengths defined by IPFIX templates. IPFIX collectors are external devices that can receive IPFIX templates and use them to interpret IPFIX logs.

EXAMPLES

```
create ipfix my_dest pool-name my_pool
```

Creates a destination named my_dest which sends IPFIX messages to the pool named my_pool.

```
delete ipfix my_dest
```

Deletes the destination named my_dest. Destinations cannot be deleted when in use by a publisher or another destination.

```
list ipfix my_dest
```

Displays properties of the destination named my_dest.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the destination belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the destination. Only the application service can modify or delete the destination.

description

A user defined description for this logging destination.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

pool-name

Specifies the LTM pool that receives messages from the IPFIX destination. This option is required for the create command. The pool should contain one or more IPFIX collectors; use the "itm pool" component to set up an LTM pool.

protocol-version

Specifies the protocol version used to encode IPFIX messages sent by this logging destination. The possible values are ipfix and netflow-9. The default is ipfix.

template-delete-delay

This feature is not implemented.

template-retransmit-interval

Specifies the time interval, in seconds, after which this IPFIX logging destination must resend all active IPFIX Templates to the pool of IPFIX collectors.

The logging destination periodically retransmits all of its IPFIX templates at the interval you set in this property. These retransmissions can be helpful if the transport-profile is UDP, a lossy transport mechanism. They can also be useful for debugging a network session with a network analyzer, such as Wireshark.

The default value is 30 seconds.

transport-profile

Specifies the name of a profile for the transport protocol to be used by this IPFIX logging destination. You can use any existing TCP-based or UDP-based profile. The default value is the default udp profile.

You can use the itm profile tcp command (see "itm profile tcp") to create a TCP profile, or itm profile udp (see "itm profile udp") to create a UDP profile.

serverssl-profile

Specifies the name of a server-side SSL profile to be used by this IPFIX Log Destination. The default is not to use a server-side SSL profile. If one is specified, the IPFIX Log Destination must be configured to use TCP as the transport protocol, and will use SSL over TCP to communicate with the configured IPFIX collectors.

You can use the itm profile server-ssl command (see "itm profile server-ssl") to create a server-side SSL profile.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

itm pool, itm profile tcp, itm profile udp, create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys log-config destination ipfix(1)

sys log-config destination local-database

NAME

local-database - Modify the Local Database destination.

MODULE

sys log-config destination

SYNTAX

Modify the Local Database component within the sys log-config destination module using the syntax shown in the following sections.

MODIFY

modify local-database [name]

options:

all
description [string]

DISPLAY

list local-database

list local-database [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line

Note: There is only one Local Database destination, local-db. This destination cannot be created or deleted.

DESCRIPTION

You can use this destination component to modify the Local Database destination for the common logging interface. There is only one Local Database destination that cannot be deleted.

EXAMPLES

list local-database local-db

Displays properties of the Local Database destination.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the modify command.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013. All rights reserved.

BIG-IP 2013-04-10 sys log-config destination local-database(1)

sys log-config destination local-syslog

NAME

local-syslog - Configures the Local Syslog destination.

MODULE

sys log-config destination

SYNTAX

Modify the Local Syslog component within the sys log-config destination module using the syntax shown in the following sections.

MODIFY

modify local-syslog [name]

options:

all
default-facility [local0 | local1 | local2 | local3 | local4 | local5 | local6 | local7]
default-severity [alert | crit | debug | emerg | err | info | notice | warn]
description [string]

DISPLAY

```
list local-syslog
list local-syslog [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  one-line
```

Note: There is only one Local Syslog destination, local-syslog. This destination cannot be created or deleted.

DESCRIPTION

You can use this destination component to modify the Local Syslog destination for the common logging interface. There is only one Local Syslog destination which cannot be deleted.

EXAMPLES

```
list local-syslog local-syslog
```

Displays properties of the Local Syslog destination.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

default-facility

Specifies the facility given to log messages received that do not already have one. The default value is local0. The options are local0, local1, local2, local3, local4, local5, local6, and local7.

default-severity

Specifies the severity given to log messages received that do not already have one. The default value is info. The options are debug, info, notice, warn, err, crit, alert, and emerg.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the modify command.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 sys log-config destination local-syslog(1)

sys log-config destination management-port

NAME

management-port - Sends received messages to a specified IP address and port through the management interface.

MODULE

sys log-config destination

SYNTAX

Configure the Management Port Log component within the sys log-config destination module using the syntax shown in the following sections.

CREATE/MODIFY

```
create management-port [name]
modify management-port [name]
options:
  all
  app-service [[string] | none]
  description [string]
  ip-address [ ip address ]
  port [ port ]
  protocol [ tcp | udp ]
```

DISPLAY

list management-port
list management-port [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line

DELETE

delete management-port [name]

Note: You must remove all references to a destination before you can delete the destination. Default destinations may not be deleted.

DESCRIPTION

You can use this destination component to create Management Port Log forwarding destinations for the common logging interface.

EXAMPLES

create management-port my_dest ip-address 1.2.3.4 port 99 protocol udp

Creates a destination named my_dest which forwards to the address 1.2.3.4:99 using the UDP protocol.

delete management-port my_dest

Deletes the destination named my_dest. Destinations cannot be deleted when in use by a publisher or another destination.

list management-port my_dest

Displays properties of the destination named my_dest.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the destination belongs. The default value is none. **Note:** If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the destination. Only the application service can modify or delete the destination.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

ip-address

Specifies the IP address that will receive messages from the specified destination.

port Specifies the port of the IP address that will receive messages from the specified destination.

protocol

Specifies the protocol used to send messages to the specified destination. The default value is tcp. The options are tcp and udp.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2015. All rights reserved.

BIG-IP 2015-07-22 sys log-config destination management-port(1)

sys log-config destination remote-high-speed-log

NAME

remote-high-speed-log - Sends received messages to a specified pool.

MODULE

sys log-config destination

SYNTAX

Configure the Remote High Speed Log component within the sys log-config destination module using the syntax shown in the following sections.

CREATE/MODIFY

create remote-high-speed-log [name]

modify remote-high-speed-log [name]

options:

all

app-service [[string] | none]

description [string]

distribution [adaptive | balanced | replicated]

pool-name [string]

protocol [tcp | udp]

DISPLAY

list remote-high-speed-log

list remote-high-speed-log [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete remote-high-speed-log [name]

Note: You must remove all references to a destination before you can delete the destination. Default destinations may not be deleted.

DESCRIPTION

You can use this destination component to create Remote High Speed Log forwarding destinations for the common logging interface.

EXAMPLES

```
create remote-high-speed-log my_dest pool-name my_pool
```

Creates a destination named my_dest which forwards to the pool my_pool.

```
delete remote-high-speed-log my_dest
```

Deletes the destination named my_dest. Destinations cannot be deleted when in use by a publisher or another destination.

```
list remote-high-speed-log my_dest
```

Displays properties of the destination named my_dest.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the destination belongs. The default value is none. **Note:** If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the destination. Only the application service can modify or delete the destination.

description

User defined description.

distribution

Specifies the distribution method used by the Remote High Speed Log destination to send messages to pool members. The default method is adaptive: connections to pool members will be added as required to provide enough logging bandwidth. This can have the undesirable effect of logs accumulating on only one pool member when it provides sufficient logging bandwidth on its own. balanced sends each successive log to a new pool member, balancing the logs among them according to the pool's load balancing method. replicated replicates each log to all pool members, for redundancy.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

pool-name

Specifies the Itm pool that receives messages from the Remote High Speed Log destination. This option is required for the create command.

protocol

Specifies the protocol used to send messages to the specified pool. The default value is tcp. The options are tcp and udp.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013. All rights reserved.

BIG-IP 2014-12-18 sys log-config destination remote-high-speed-log(1)

sys log-config destination remote-syslog

NAME

remote-syslog - Configures Remote Syslog destinations to format log messages into Syslog format and forward them to a Remote High-Speed Log destination.

MODULE

sys log-config destination

SYNTAX

Configure the Remote Syslog component within the sys log-config destination module using the syntax shown in the following sections.

CREATE/MODIFY

create remote-syslog [name]

modify remote-syslog [name]

options:

all

app-service [[string] | none]

default-facility [local0 | local1 | local2 | local3 | local4 | local5 | local6 | local7]

default-severity [alert | crit | debug | emerg | err | info | notice | warn]

description [string]

format [legacy-bigip | rfc3164 | rfc5424]

remote-high-speed-log [string]

DISPLAY

list remote-syslog

list remote-syslog [[[name] | [glob] | [regex]] ...]

options:

all-properties

non-default-properties

one-line

DELETE

delete remote-syslog [name]

Note: You must remove all references to a destination before you can delete the destination. Default destinations may not be deleted.

DESCRIPTION

You can use this destination component to create Remote Syslog formatting destinations for the common logging interface.

EXAMPLES

```
create remote-syslog my_dest remote-high-speed-log another_dest
```

Creates a destination named my_dest which forwards to another destination another_dest. another_dest may not be another Remote Syslog destination.

```
delete remote-syslog my_dest
```

Deletes the destination named my_dest. Destinations cannot be deleted when in use by a publisher or another destination.

```
list remote-syslog my_dest
```

Displays properties of the destination named my_dest.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the destination belongs. The default value is

none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the destination. Only the application service can modify or delete the destination.

default-facility
Specifies the facility given to log messages received that do not already have a facility listed. The default value is local0. The options are local0, local1, local2, local3, local4, local5, local6, and local7.

default-severity
Specifies the severity given to log messages received that do not already have a severity listed. The default value is info. The options are debug, info, notice, warn, err, crit, alert, and emerg.

description
User defined description.

format
Specifies the syslog format received messages are formatted into. The default value is rfc3164. The options are legacy-bigip, rfc3164, and rfc5424. For more information, see the respective RFCs.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

remote-high-speed-log
Specifies the forwarding destination to send logs in the syslog format. This option is required for the create command. It may only be a remote high speed log destination or a management port destination.

SEE ALSO
create, delete, glob, list, modify, regex, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013. All rights reserved.

BIG-IP 2014-12-18 sys log-config destination remote-syslog(1)

sys log-config destination splunk

NAME
splunk - Configures Splunk formatting destinations to format incoming log messages into the Splunk format.

MODULE
sys log-config destination

SYNTAX
Configure the Splunk component within the sys log-config destination module using the syntax shown in the following sections.

CREATE/MODIFY
create splunk [name]
modify splunk [name]
options:
all
app-service [[string] | none]
description [string]
forward-to [string]

DISPLAY
list splunk
list splunk [[[name] | [glob] | [regex]] ...]
options:
all-properties
non-default-properties
one-line

DELETE
delete splunk [name]

Note: You must remove all references to a destination before you can delete the destination. Default destinations may not be deleted.

DESCRIPTION

You can use this destination component to create Splunk formatting destinations for the common logging interface.

EXAMPLES

```
create splunk my_dest forward-to another_dest
```

Creates a destination named `my_dest` which forwards to another destination `another_dest`. `another_dest` must be a Local Syslog, Local Database, Remote Syslog, or Remote High Speed Log destination.

```
delete splunk my_dest
```

Deletes the destination named `my_dest`.

```
list splunk my_dest
```

Displays properties of the destination named `my_dest`. Destinations cannot be deleted when in use by a publisher.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the destination belongs. The default value is none. **Note:** If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the destination. Only the application service can modify or delete the destination.

description

User defined description.

forward-to

Specifies a Local Syslog, Local Database, Remote Syslog, or Remote High Speed Log destination to receive Splunk formatted log messages. This is required for the creation of a Splunk destination.

glob Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands `create`, `delete`, and `modify`.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

SEE ALSO

`create`, `delete`, `glob`, `list`, `modify`, `regex`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 sys log-config destination splunk(1)

sys log-config filter

NAME

`filter` - Configures the filter component which filters out log messages for distribution by the publisher component.

MODULE

`sys log-config`

SYNTAX

Configure the filter component within the `sys log-config` module using the syntax shown in the following sections.

CREATE/MODIFY

```
create filter [name]
```

```
modify filter [name]
```

options:

`all`

app-service [[string] | none]
description [string]
level [alert | crit | debug | emerg | err | info | notice | warn]
message-id [8 digit hex number | none]
publisher [[string] | none]
source [accesscontrol | accessperrequest | adapt | adfs-proxy | alertd | all |
api-protection | apmac | arp | authz | autodisc | autodsd | avr |
based | bcm56xxd | bdosd | big3d | big3dshim | bigd | bigdb | bigdbd | bigpipe |
bigstart | bp | keymgmt | checkcert | chmand | cifs | clusterd | coapi | common |
common-f5logging | common-fpdd | config-db | connapi | cs | cssd | csyncd |
daemon | debugd | deflate | devmgmt | diameter | dmon | dosprotect | dummy | dwbl | dynad |
eca | em-admin | em-alert | em-clientlib | em-common | em-device | em-discovery |
em-file | em-lib | em-report | em-stats | em-swim | errdefsd | eventd | evrouted |
fflag | fips | firewall-FQDN | firewall-nat | fix | ftp | get-dossier | gtmd | gtp |
guestagentd | ha | ha-table | halmsg | hclientd | hornet-lib | hornet-nest |
hornet-nest-flow-manager | hornet-nest-updater | hornet-neuron-updater | hornet-server |
hornet-text-client | hostagentd | http | htconnector | hwctl | hwpd |
icrd | imap | ip | ipfix | ipfix-proxy | ipfixrules | iprepd | ipsec | isession | istatsd |
ivs | lacpd | layer4 | libhal | lind | lldpd | localdb | lopd | lsn | lsnapi | mamidbridged |
map | mapi | mcp | mcpd | mcpd-apm | mcpd-asm | mcpd-centmgmt | mcpd-clustering |
mcpd-dev | mcpd-dpi | mcpd-firewall | mcpd-framework | mcpd-gtm | mcpd-ips | mcpd-ltm |
mcpd-net | mcpd-pem | mcpd-sys | mcpd-wam | mcpd-woc |
mdm | mgmt-acld | mr | mrsip | msgbusd | mysqlhad | natstatsd | net |
network | no-source | packet-filter | pccd | pcp | pem | pfmand | pgadmind |
pkcs11d | pktclass | plugin | policy | pop3 | portal-access | pptp | probe-plusplus |
promptstatusd | provisioning | pva | pvad | qkcloud | radius | ramcache | rba |
rewrite | rtsp | rules |
saspd | scim | scriptd | sctp | sdmd | sflow | shell | shmapd | smtps | snmp | sod |
spolicy | ssl | ssl-orchestrator | sso | stated | statsd | statusd | stmm | stpd | subagents | swg |
syscall | system-check | tamd | tcl-checker | tcpdump | tftp | tmm | tmm-tcp |
tmrouted | tmsd | ts | tunnel | urld | urldb | urldbmgd | vcmpd | vdi | vxland |
webssh | websso | woc-plugin | wr-urldb | xconfig | xdb | zfd | zxfrd | gpa | cryptod |
icr-eventd | ips | dpi]

DISPLAY

list filter

list filter [[[name] | [glob] | [regex]] ...]

options:

- all-properties
- non-default-properties
- one-line

DELETE

delete filter [name]

DESCRIPTION

You can use the filter component to configure the filters for the common logging interface.

EXAMPLES

```
create filter my_filter publisher my_pub
```

Creates a filter named my_filter with the publisher my_pub.

```
delete filter my_filter
```

Deletes the filter named my_filter.

```
list filter my_filter
```

Displays properties of the filter named my_filter.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the filter belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the filter. Only the application service can modify or delete the filter.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

level

The minimum severity level of logs to be filtered. The severity levels in increasing order are debug, info, notice, warn, err, crit, alert, and emerg. The default value is debug.

message-id

A refinement for filtering out specific logs. The default value is none. This is an eight digit hex number. The proper hex value can be obtained from an existing log message by extracting the eight digit value.

For example, the message-id for the example log message below is highlighted.

```
Oct 9 15:38:20 bigip1 notice mcpd[21498]: 01070410:5: Removed subscription with subscriber id logstatd
```

name Specifies a unique name for the component. This option is required for the commands **create**, **delete**, and **modify**.

publisher
A publisher to send filtered log messages. The default value is **none**.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (**@**[regular expression]) to indicate that the identifier is a regular expression. See help **regex** for a description of regular expression syntax.

source
The stream of log messages that will be filtered by the created/modified filter. The default value is **all**.

SEE ALSO

create, **delete**, **glob**, **list**, **modify**, **regex**, **tmsh**

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013, 2015. All rights reserved.

BIG-IP 2019-09-11 **sys log-config filter(1)**

sys log-config publisher

NAME

publisher - Configures lists of destinations for the common logging interface.

MODULE

sys log-config

SYNTAX

Configure the publisher component within the **sys log-config** module using the syntax shown in the following sections.

CREATE/MODIFY

```
create publisher [name]
modify publisher [name]
options:
all
app-service [[string] | none]
description [string]
destinations [add | delete | none | replace-all-with] {
  [ [destinations] ]
}
```

DISPLAY

```
list publisher
list publisher [ [ [name] | [glob] | [regex] ] ... ]
options:
all-properties
non-default-properties
one-line
```

DELETE

```
delete publisher [name]
```

Note: You must remove all references to a publisher before you can delete the publisher. Default publishers may not be deleted.

DESCRIPTION

You can use the publisher component to configure publishers for the common logging interface.

EXAMPLES

```
create publisher my_pub destinations add {
destination_1
destination_2
}
```

Creates a publisher named **my_pub** with two destinations, **destination_1** and **destination_2**.

```
delete publisher my_pub
```

Deletes the publisher named **my_pub**.

list publisher my_pub

Displays properties of the publisher named my_pub.

OPTIONS

all Specifies that you want to modify all of the existing components of the specified type.

app-service

Specifies the name of the application service to which the publisher belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the publisher. Only the application service can modify or delete the publisher.

description

User defined description.

destinations

Adds, deletes, or replaces a set of destinations.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2013. All rights reserved.

BIG-IP 2013-04-10 sys log-config publisher(1)

sys log-rotate

NAME

log-rotate - Configures log rotation for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the log-rotate component within the sys module using the syntax in the following sections.

MODIFY

modify log-rotate

options:

common-backlogs [integer]

common-include [string]

description [string]

ilx-include [string]

ilx-rotations [string]

ilx-schedule [string]

ilx-size [string]

include [string]

max-file-size [integer]

mysql-include [string]

syslog-include [string]

tomcat-include [string]

wa-include [string]

edit log-rotate

options:

all-properties

non-default-properties

DISPLAY

list log-rotate

list log-rotate [option]

show running-config log-rotate

show running-config log-rotate [option]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can configure the system to rotate the log files after a specified length of time. This helps to clear the hard drive of unneeded log files.

EXAMPLES

modify log-rotate common-backlogs 7

Specifies that the system saves seven copies of the common log files.

list log-rotate all-properties

Displays the configuration of the log-rotate component.

OPTIONS

common-backlogs

Specifies the number of logs that you want the system to save. Select a number from the valid range of 1 - 100. The default value is 24.

common-include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the common-include option incorrectly, you put the functionality of the system at risk.

description

User defined description.

ilx-include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the include option incorrectly, you put the functionality of the system at risk.

ilx-rotations

The default value is 10.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the include option incorrectly, you put the functionality of the system at risk.

ilx-schedule

The default value is daily.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the include option incorrectly, you put the functionality of the system at risk.

ilx-size

The value is specified in kilobytes. The default value is 10240 kilobytes.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the include option incorrectly, you put the functionality of the system at risk.

include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the include option incorrectly, you put the functionality of the system at risk.

max-file-size

The max size of rotated log files in kB. The default value is 1024000.

syslog-include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the syslog-include option incorrectly, you put the functionality of the system at risk.

tomcat-include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the tomcat-include option incorrectly, you put the functionality of the system at risk.

wa-include

The default value is none.

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using this option. If you use the wa-include option incorrectly, you put the functionality of the system at risk.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2016-07-29 sys log-rotate(1)

sys log

NAME

log - Displays various system log files.

MODULE

sys

SYNTAX

Configure the log component within the sys module using the syntax in the following sections.

DISPLAY

show log

show log [audit | daemon | gtm | kernel | ltm | mail | messages | security | tmm | user | webui]

options:

lines [integer]

range [date range]

DESCRIPTION

You can use the log component to display various logs.

EXAMPLES

show log

Displays a list of logs that you can view.

show log gtm

Displays the Global Traffic Manager log.

show log gtm lines 100 range 2/19/2006:15:04:00--epoch

Displays no more than 100 lines of the Global Traffic Manager log that were logged before the 19th of February 2006 at 3:04 pm.

OPTIONS

audit

Displays a log of configuration changes.

daemon

Displays the Unix daemon logs.

gtm Displays the Global Traffic Manager logs.

kernel

Displays Linux Kernel messages.

lines

Specifies how many lines of the log that you want the system to display at one time.

ltm Displays Local Traffic Manager logs.

mail Displays mail daemon logs.

messages

Displays application messages.

range

Specifies the date range of the log information that you want the system to display.

security
Displays security-related messages.

tmm Displays Traffic Manager Micro-kernel logs.

user Displays various user process logs.

webui
Displays Configuration utility logs.

SEE ALSO
show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012-2013. All rights reserved.

BIG-IP 2013-04-12 sys log(1)

sys mac-address

NAME
mac-address - Displays all MAC addresses currently associated with a configuration object on a BIG-IP(r) system, including all dynamically-discovered MAC addresses.

MODULE
sys

SYNTAX
Display the MAC addresses associated with a BIG-IP system configuration using the syntax in the following section.

DISPLAY
show mac-address
options:
field-fmt

DESCRIPTION
You can use the mac-address component to display the location on the BIG-IP system of the MAC addresses associated with a configuration object. The system displays the following information, which identifies the location of the MAC address in the configuration.

Entry
Displays the MAC address.

Component
Displays the type of component associated with the MAC address, for example, net interface.

Object-ID
Displays the name of a configuration object associated with the MAC address, for example, 2.1.

Property
Displays the name of the property that contains the MAC address value. Note that if the MAC address is an object identifier the system displays n/a.

EXAMPLES
show mac-address

Displays all MAC addresses currently associated with a BIG-IP system configuration object.

OPTIONS
For information about the options that you can use with the command show, see help show.

SEE ALSO
show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-11 sys mac-address(1)

sys management-dhcp

NAME

management-dhcp - Configures dhcp settings for the management interface (MGMT). The changes in this object are reflected in dhclient's next lease renewal cycle and doesn't effect the current lease. User can force the changes to take effect right away by restarting dhclient.

MODULE

sys

SYNTAX

Configure the management-dhcp component within the sys module using the syntax in the following sections.

MODIFY

```
modify management-dhcp [name]
options:
  client-id [string]
  description [string]
  hostname [string]
  request-options [add | delete | modify | replace-all-with]
  send-options [add | delete | modify | replace-all-with]
  supersede-options
    [add | delete | modify | replace-all-with] {
      [name] ... {
        value
          [add | delete | modify | replace-all-with] {
            }
          }
    }
  supersede-options none
```

```
edit management-dhcp [name]
```

```
options:
  all-properties
```

DISPLAY

```
list management-dhcp
list management-dhcp [name]
show running-config sys management-dhcp
show running-config sys management-dhcp [name]
options:
  all-properties
  one-line
```

DESCRIPTION

Specifies DHCP client settings for the management interface. These settings will be used to retrieve an IP address for the management interface if mgmt-dhcp is enabled.

EXAMPLES

```
modify management-dhcp default request-options add ntp-servers
```

Adds ntp-servers to the lists of options requested by the management interface DHCP client.

OPTIONS

client-id

Specifies the client identifier to send to the DHCP server.

description

User defined description.

hostname

Specifies the hostname to send to the DHCP server.

request-options

Specifies the options to request from the DHCP server. Adding or removing an option will be reflected at next lease renewal with dhcp server or upon restarting DHCP client.

send-options

Specifies the options to send to the DHCP server.

supersede-options

Specifies dhclient options for which BIG-IP should always use a locally-configured value or values rather than whatever is provided by the DHCP server in the lease.

SEE ALSO

edit, list, modify, show, sys management-ip, sys management-route, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

BIG-IP 2018-01-11 sys management-dhcp(1)

sys management-ip

NAME

management-ip - Configures the ip address and netmask for the management interface (MGMT).

MODULE

sys

SYNTAX

Configure the management-ip component within the sys module using the syntax in the following sections.

CREATE/MODIFY

```
create management-ip [ip address/netmask]
create management-ip [ip address/prefixlen]
modify management-ip [ip address/prefixlen]
options:
description
```

DISPLAY

```
list management-ip
show running-config management-ip
options:
all-properties
one-line
```

DELETE

```
delete management-ip [ip address/netmask]
delete management-ip [ip address/prefixlen]
```

DESCRIPTION

Specifies network settings for the management interface.

The management interface is available on all switch platforms and is designed for management purposes. You can access the browser-based Configuration utility and command line configuration utility through the management port. You cannot use the management interface in traffic management VLANs. You can configure only one IP address on the management interface.

After you make any changes using the management-ip component, issue the following command sequence to save the changes to the bigip_base.conf file: save sys config.

To configure management-ip firewall rules, see security firewall management-ip-rules.

Note: modify only allows modification of the description field. If you wish to change the IP address of the management interface, please see the example below.

EXAMPLES

```
create management-ip 10.2.3.4/255.255.0.0
```

Creates the IP address 10.2.3.4 on the management interface.

```
create management-ip 10.2.3.4/16
```

Creates the IP address 10.2.3.4 on the management interface.

```
delete sys management-ip 10.2.3.4/25; create sys management-ip 10.2.3.5/25
```

Changes the IP address of the management interface. Note: modify does not allow a user to change the IP address directly.

OPTIONS

[ip address/netmask]

Specifies the IPv4 address and netmask. DHCP should be set to 'disabled' or 'dhcpv6' before creating IPv4 management-ip.

[ip address/prefixlen]

Specifies the IPv6 address and prefix length. DHCP should be set to 'disabled' or 'dhcpv4' before creating IPv6 management-ip.

description

User defined description.

dhcp-enabled

Specifies if the ip address has been configured by DHCP.

SEE ALSO

create, delete, list, modify, save, show, security firewall management-ip-rules, sys management-route, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2012. All rights reserved.

BIG-IP 2018-06-14 sys management-ip(1)

sys management-ovsdb

NAME

management-ovsdb - Configures the OVSDB server.

MODULE

sys

SYNTAX

Configure the management-ovsdb component within the sys module using the syntax in the following sections.

MODIFY

modify management-ovsdb

options:

bfd-disabled

bfd-enabled

bfd-route-domain [route-domain]

ca-cert-file [filename]

cert-file [filename]

cert-key-file [filename]

controller-addresses [IP address(es)]

description [string]

disabled

enabled

flooding-type [multipoint | replicator]

log-level [level]

logical-routing-type [none | backhaul]

port [port number]

tunnel-floating-addresses [IP address(es)]

tunnel-local-address [IP address]

tunnel-maintenance-mode [active | passive]

DISPLAY

list management-ovsdb

show running-config management-ovsdb

options:

all-properties

one-line

bfd-disabled

bfd-enabled

bfd-route-domain

ca-cert-file

cert-file

cert-key-file

controller-addresses

description

disabled

enabled

flooding-type

log-level

logical-routing-type

port

tunnel-floating-addresses

tunnel-local-address

tunnel-maintenance-mode

DESCRIPTION

Specifies the configurations for the OVSDB server.

EXAMPLES

modify management-ovsdb controller-addresses add { 10.0.0.1 }

Specifies 10.0.0.1 as the controller address.

list management-ovsdb all-properties

Displays the OVSDB server configurations.

OPTIONS

bfd-disabled

Disable the BFD sessions between the BIG-IP and replicators.

bfd-enabled

Enable the BFD sessions between the BIG-IP and replicators.

bfd-route-domain

Specifies the route-domain on which the VXLAN tunnel of VNI 0 is created for the BFD sessions between the BIG-IP and replicators.

ca-cert-file

Specifies the name of the CA certificate file.

cert-file

Specifies the name of the certificate file.

cert-key-file

Specifies the name of the certificate key file.

controller-addresses

Specifies the IP address(es) of the controller.

disabled

Disables OVSDDB management.

enabled

Enables OVSDDB management.

flooding-type

Specifies the flooding type to use to transmit unknown destination frames.

log-level

Specifies the log level for OVSDDB management. The log file is located at `/var/log/vxland.log`.

logical-routing-type

Specifies the logical routing type.

port

Specifies the OVSDDB connection port. The default port is 6640.

tunnel-floating-addresses

Specifies the floating endpoint address(es) for the tunnels. The addresses need to be a self IP address associated with a floating traffic-group.

tunnel-local-address

Specifies the local endpoint address for the tunnels created by the controller. A valid IP address for the tunnel local endpoint is required when OVSDDB management is enabled.

tunnel-maintenance-mode

Specifies whether the tunnels are created automatically (active) or manually (passive).

description

User defined description.

SEE ALSO

`list`, `modify`, `save`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2015. All rights reserved.

BIG-IP 2017-07-17 `sys management-ovsdb(1)`

sys management-proxy-config

NAME

`management-proxy-config` - Configures proxy configuration for database download.

MODULE

`sys`

SYNTAX

Configure a `management-proxy-config` component within the `sys` module using the syntax shown in the following sections.

CREATE/MODIFY

create management-proxy-config [name]
modify management-proxy-config [name]
options:
 description [string]
 proxy-ip-addr [ip address]
 proxy-port [port]
 username [string]
 password [string]

edit management-proxy-config [name]
options:
 all-properties
 non-default-properties

DISPLAY

list management-proxy-config [name]
options:
 all-properties
 non-default-properties
 one-line

DELETE

delete management-proxy-config [name]

DESCRIPTION

Configures proxy configuration for database download. The management-proxy-config consists of the object name, proxy ip address (proxy-ip), proxy port (proxy-port), username and password. You can have only one proxy configuration specified.

EXAMPLES

```
create management-proxy-config test-proxy { proxy-ip-addr 10.10.10.10 proxy-port 1010 username test password test }
```

Creates the proxy configuration with name test-proxy and uses the specified proxy-ip-addr 10.10.10.10 and proxy-port 1010 with supplied credentials username/password test/test for database download.

```
modify management-proxy-config test-proxy { proxy-port 1012 }
```

Modify the proxy configuration for database download to use proxy port 1012.

```
delete management-proxy-config test-proxy
```

Delete the test-proxy configuration for database download.

OPTIONS

description
User defined description.

proxy-ip-addr
Specifies proxy IP to be used for database download.

proxy-port
Specifies proxy port to be used for database download (default value is 3128).

username
Specifies username for proxy configuration to be used for database download.

password
Specifies password for proxy configuration to be used for database download.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013-2014, 2016. All rights reserved.

BIG-IP 2017-01-26 sys management-proxy-config(1)

sys management-route

NAME

management-route - Configures route settings for the management interface (MGMT).

MODULE

sys

SYNTAX

Configure the management-route component within the sys module using the syntax in the following sections.

CREATE/MODIFY

create management-route [name | default | default-inet6]

options:

description [string]
gateway [ip address]
mtu [number]
network [ip address/netmask]
type [interface | blackhole]

modify management-route [name | default | default-inet6]

options:

description [string]
gateway [ip address]
mtu [number]
type [interface | blackhole]

edit management-route [[name | default | default-inet6]
| [glob] | [regex]] ...]

options:

all-properties

DISPLAY

list management-route

list management-route [[name | default | default-inet6]
| [glob] | [regex]] ...]

show running-config management-route

show running-config management-route [[name | default
| default-inet6] | [glob] | [regex]] ...]

options:

all-properties
one-line

DELETE

delete management-route [name]

DESCRIPTION

Specifies route settings for the management interface. You must configure a route on the management interface if you want to access the management network on the BIG-IP(r) system by connecting from another network.

The management interface is available on all switch platforms and is designed for management purposes. You can access the browser-based Configuration utility and command line configuration utility through the management port. You cannot use the management interface in traffic management VLANs.

EXAMPLES

create management-route default gateway 10.10.10.254

Sets the management interface default gateway IP address to 10.10.10.254.

create management-route myMgmtRoute network 10.10.10.0/24 gateway 10.10.10.254

Creates a management route named myMgmtRoute for the subnet 10.10.10.0/24 whose gateway IP address is 10.10.10.254.

modify management-route 10.10.10.0/24 gateway 172.24.74.62

Changes the management interface to subnet 10.10.10.0/24, and the gateway to 172.24.74.62.

OPTIONS

default

Specifies that the system forwards packets to the destination through the default IP address and netmask, 0.0.0.0 0.0.0.0.

default-inet6

Specifies that the system forwards packets to the destination through the default version 6.0 IP address and netmask.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

[ip address/netmask]

Specifies the IP address and netmask through which the system forwards packets to the destination. You can use either of these formats: 0.0.0.0/0 or 0.0.0.0 0.0.0.0.

gateway

Specifies that the system forwards packets to the destination through the gateway with the specified IP address.

mtu Specifies the maximum transmission unit (MTU) for the management interface. The value of the MTU is the largest size that the BIG-IP system allows for an IP datagram passing through the management interface.

network

The subnet and netmask to be used for the route. This is an optional field; if empty the name should be

of the form [ip address/netmask].

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

type Specifies that traffic should be delivered to the management interface (interface) or be dropped by the system (blackhole).

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, sys management-ip, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013. All rights reserved.

BIG-IP 2015-09-24 sys management-route(1)

sys mcp-state

NAME

mcp-state - Displays information about the mcpd daemon.

MODULE

sys

SYNTAX

Display information about the mcpd daemon using mcp-state component within the sys module using the syntax in the following section.

DISPLAY

show mcp-state
options:
field-fmt

DESCRIPTION

You can use the mcp-state component to display the current state of the mcpd daemon.

EXAMPLES

show mcp-state

Displays, in a table, information about the state of the mcpd daemon.

show mcp-state field-fmt

Displays, in field format, information about the state of the mcpd daemon.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-11 sys mcp-state(1)

sys memory

NAME

memory - Displays system memory information and statistics.

MODULE

sys

SYNTAX

Configure the memory component within the sys module using the syntax in the following sections.

DISPLAY

show memory

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

global

DESCRIPTION

You can use the memory component to display information about the system memory.

EXAMPLES

show memory gig

Displays memory statistics in gigabytes.

show memory raw

Displays raw memory statistics.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 sys memory(1)

sys nethsm async-queue-stat

NAME

async-queue-stat - Display or reset Network HSM pkcs11d daemon async queue statistics

MODULE

sys

SYNTAX

Display and reset the async-queue-stat component within the sys nethsm module using the syntax in the following section.

MODIFY

reset-stats async-queue-stat

DISPLAY

show async-queue-stat

options:

(default | field-fmt)

DESCRIPTION

You can use the async-queue-stat component to display Network HSM pkcs11d async-queue statistics like queued operations and queue times. You can also reset the Network HSM pkcs11d async-queue statistics to zero at any time.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013, 2016. All rights reserved.

BIG-IP 2017-01-20 sys nethsm async-queue-stat(1)

sys nethsm pkcs11d-stat

NAME

pkcs11d-stat - Display or reset Network HSM pkcs11d daemon statistics

MODULE

sys

SYNTAX

Display and reset the pkcs11d-stat component within the sys nethsm module using the syntax in the following section.

MODIFY

reset-stats pkcs11d-stat

DISPLAY

show pkcs11d-stat

options:

(default | field-fmt)

DESCRIPTION

You can use the pkcs11d-stat component to display Network HSM pkcs11d statistics like sign, decrypt and find-key counts. You can also reset the Network HSM pkcs11d statistics to zero at any time.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013, 2016. All rights reserved.

BIG-IP 2017-01-20 sys nethsm pkcs11d-stat(1)

sys nethsm sync-queue-stat

NAME

sync-queue-stat - Display or reset Network HSM pkcs11d daemon sync queue statistics

MODULE

sys

SYNTAX

Display and reset the sync-queue-stat component within the sys nethsm module using the syntax in the following section.

MODIFY

reset-stats sync-queue-stat

DISPLAY

show sync-queue-stat

options:

(default | field-fmt)

DESCRIPTION

You can use the sync-queue-stat component to display Network HSM pkcs11d sync-queue statistics like queued operations and queue times. You can also reset the Network HSM pkcs11d sync-queue statistics to zero at any time.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011, 2013, 2016. All rights reserved.

BIG-IP 2017-01-20 sys nethsm sync-queue-stat(1)

sys ntp

NAME

ntp - Configures the Network Time Protocol (NTP) daemon for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the ntp component within the sys module using the following syntax.

MODIFY

modify ntp

options:

```
description [string]
include [string]
restrict [add | delete | replace-all-with] {
  [string] {
    address [IP address]
    default-entry [enabled | disable]
    description [string]
    ignore [enabled | disable]
    kod [enabled | disable]
    limited [enabled | disable]
    low-priority-trap [enabled | disable]
    mask [IP address]
    no-modify [enabled | disable]
    non-ntp-port [enabled | disable]
    no-peer [enabled | disable]
    no-query [enabled | disable]
    no-serve-packets [enabled | disable]
    no-trap [enabled | disable]
    no-trust [enabled | disable]
    ntp-port [enabled | disable]
    version [enabled | disable]
  }
}
restrict none
servers [add | delete | replace-all-with] {
  [hostname | IP address] ...
}
servers none
timezone [string]
```

edit ntp

options:

```
all-properties
non-default-properties
```

DISPLAY

list ntp

list ntp [option]

show running-config ntp

show running-config ntp [option]

options:

```
all-properties
non-default-properties
one-line
```

DESCRIPTION

You can use this component to configure the NTP servers for the system.

EXAMPLES

```
modify ntp servers add {192.168.1.245}
```

Adds the NTP server with the IP address, 192.168.1.245, to the system.

```
modify ntp servers replace-all-with {time.f5net.com}
```

Replaces the existing list of NTP servers with a single host, time.f5net.com.

```
modify ntp timezone "America/Los_Angeles"
```

Sets the system time to Pacific Standard Time.

```
modify ntp restrict add { basicrestrict { default-entry enable ignore enable } }
```

Adds a default restriction denying all packets.

OPTIONS

description

User defined description.

include

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using the include option. If you use this option incorrectly, you put the functionality of the system at risk.

restrict

Specifies a set of access restrictions.

address

The address for the entry. See also, the mask option. The default value is 0.0.0.0.

default-entry

Specifies whether the entry is the default entry. The default value is disabled.

description

User defined description.

ignore

Specifies whether all packets will be ignored. The default value is disabled.

kod Specifies whether a kod (kiss of death) packet will be sent when an access violation occurs. The default value is disabled.

limited

Specifies whether service will be denied if packet spacing limits are violated. The default value is disabled.

low-priority-trap

Specifies whether lower priority traps will be overridden by normal priority traps. The default value is disabled.

mask The mask for the entry. See also, the address option. The default value is 0.0.0.0.

no-modify

Specifies whether ntpq and ntpdc queries that attempt to modify the server are allowed. The default value is disabled.

non-ntp-port

When enabled, the restrict entry will be matched only if the source port is not the standard NTP UDP port (123). The default value is disabled.

no-peer

Specifies whether packets will be denied if they mobilize a new association. The default value is disabled.

no-query

Specifies whether ntpq and ntpdc queries will be denied. The default value is disabled.

no-serve-packets

Specifies whether all queries except ntpq and ntpdc will be denied. The default value is disabled.

no-trap

Specifies whether to decline the mode 6 control message trap service to matching hosts. The default value is disabled.

no-trust

Specifies whether to reject packets that are not cryptographically authenticated. The default value is disabled.

ntp-port

When enabled, the restrict entry will be matched only if the source port is the standard NTP UDP port (123). The default value is disabled.

version

Specifies whether packets will be rejected if they do not match the local NTP version. The default values is disabled.

servers

Configures NTP servers for the BIG-IP system.

timezone

Specifies the time zone that you want to use for the system time.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-10-06 sys ntp(1)

sys outbound-smtp

NAME

outbound-smtp - Configures outgoing email for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the outbound-smtp component within the sys module using the following syntax.

MODIFY

modify outbound-smtp

options:

description [string]

from-line-override [string]

mailhub [string]

rewrite-domain [string]

edit outbound-smtp

options:

all-properties

non-default-properties

DISPLAY

list outbound-smtp

list outbound-smtp [option]

show running-config outbound-smtp

show running-config outbound-smtp [option]

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use this component to configure the outgoing SMTP server that the system will use to send automated email.

EXAMPLES

modify outbound-smtp mailhub smtp.yoursite.com:587

Configures the TMOS system to send outgoing email through the specified SMTP server.

OPTIONS

description

User defined description.

from-line-override

Specify whether From field's domain can be overridden.

mailhub

The SMTP server to use to send outgoing automated email.

rewrite-domain

The domain name in the From address to use when sending outgoing automated email.

SEE ALSO

edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2013. All rights reserved.

sys performance all-stats

NAME

all-stats - Resets or displays all performance statistics.

MODULE

sys performance

SYNTAX

Reset or display all performance statistics for the system within the sys_performance module using the syntax in the following sections. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

MODIFY

reset-stats all-stats

DISPLAY

show all-stats

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
(detail | historical)

DESCRIPTION

You can use the all-stats component to reset or display all system performance statistics.

Note that tmsh only displays performance statistics when you explicitly request them.

EXAMPLES

show all-stats detail

Displays detailed information about system performance in the system default units.

reset-stats all-stats

Resets all performance statistics for the system.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, sys performance connections, sys performance gtm, sys performance ramcache, sys performance system, sys performance throughput, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012. All rights reserved.

sys performance connections

NAME

connections - Displays connection performance information.

MODULE

sys performance

SYNTAX

Display statistics for the connections component within the sys_performance module using the syntax in the following section. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

DISPLAY

show connections

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
(detail | historical)

DESCRIPTION

You can use the connections component to display information about system performance, including details about new and active connections and HTTP requests.

You can reset the connection performance statistics using the all-stats component.

EXAMPLES

show connections gig detail

Displays detailed information about connection performance in gigabytes.

show connections historical

Displays historical performance information about connections.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys performance all-stats, sys performance gtm, sys performance ramcache, sys performance system, sys performance throughput, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2012-03-26 sys performance connections(1)

sys performance dnsexpress

NAME

dnsexpress - Displays performance information for the DNS-Express.

MODULE

sys performance

SYNTAX

Display statistics for the dnsexpress component within the sys performance module using the syntax in the following section. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

DISPLAY

show dnsexpress

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
(historical)

DESCRIPTION

You can use the dnsexpress component to display information about system performance, including the number of queries, responses, zone transfer messages, and NOTIFYs.

EXAMPLES

show dnsexpress historical

Displays historical performance information for DNS-Express.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys performance all-stats, sys performance connections, sys performance ramcache, sys performance system, sys performance throughput, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-10-22 sys performance dnsexpress(1)

sys performance dnssec

NAME

dnssec - Displays performance information for the DNSSEC signing.

MODULE

sys performance

SYNTAX

Display statistics for the dnssec component within the sys performance module using the syntax in the following section. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

DISPLAY

show dnssec

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
(historical)

DESCRIPTION

You can use the dnssec component to display information about system performance, including the number of queries for specific DNSSEC types and zone transfer signing.

EXAMPLES

show dnssec historical

Displays historical performance information for DNSSEC signing.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys performance all-stats, sys performance connections, sys performance ramcache, sys performance system, sys performance throughput, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2014. All rights reserved.

BIG-IP 2014-10-22 sys performance dnssec(1)

sys performance gtm

NAME

gtm - Displays performance information for the Global Traffic Manager.

MODULE

sys performance

SYNTAX

Display statistics for the gtm component within the sys performance module using the syntax in the following section. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

DISPLAY

show gtm

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
(detail | historical)

DESCRIPTION

You can use the gtm component to display information about system performance, including details about the Global Traffic Manager, including number of requests, resolutions, persisted connections, and those returned to DNS.

You can reset the Global Traffic Manager performance statistics using the all-stats component.

EXAMPLES

show gtm detail

Displays detailed performance information about the Global Traffic Manager in the system default units.

show gtm historical

Displays historical performance information about the Global Traffic Manager.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys performance all-stats, sys performance connections, sys performance ramcache, sys performance system, sys performance throughput, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2012-03-26 sys performance gtm(1)

sys performance ramcache

NAME

ramcache - Displays RAM cache performance information.

MODULE

sys performance

SYNTAX

Display statistics for the ramcache component within the sys performance module using the syntax in the following section. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

DISPLAY

show ramcache

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

historical

DESCRIPTION

You can use the ramcache component to display RAM cache utilization information.

You can reset the RAM cache performance statistics using the all-stats component.

EXAMPLES

show ramcache default

Displays ramcache performance information in the system default units.

show ramcache historical

Displays historical ramcache performance information.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys performance all-stats, sys performance connections, sys performance gtm, sys performance system, sys performance throughput, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2012-03-26 sys performance ramcache(1)

sys performance system

NAME

system - Displays system performance information.

MODULE

sys performance

SYNTAX

Display statistics for the system component within the sys performance module using the syntax in the following section. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

DISPLAY

show system

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

(detail | historical)

DESCRIPTION

You can use the system component to display CPU and memory usage information.

You can reset the system performance statistics using the all-stats component.

EXAMPLES

show system detail

Displays detailed system performance information in the system default units.

show system historical

Displays historical system performance information.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys performance all-stats, sys performance connections, sys performance gtm, sys performance ramcache, sys performance throughput, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2012-03-26 sys performance system(1)

sys performance throughput

NAME

throughput - Displays performance information about traffic throughput.

MODULE

sys performance

SYNTAX

Display statistics for the throughput component within the sys performance module using the syntax in the following section. On VIPRION(r) systems, displaying performance statistics on a secondary blade is not supported.

DISPLAY

show throughput

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

(detail | historical)

DESCRIPTION

You can use the throughput component to display information about traffic throughput, including client, server, compression, and SSL transactions.

You can reset the throughput performance statistics using the all-stats component.

EXAMPLES

show throughput gig detail

Displays detailed throughput performance information in gigabits per second.

show throughput historical

Displays historical throughput performance information.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys performance all-stats, sys performance connections, sys performance gtm, sys performance ramcache, sys performance system, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2012-03-26 sys performance throughput(1)

sys pfman consumer

NAME

consumer - Manage the state of pfman health status consumers.

MODULE

sys pfman

SYNTAX

Manage pfman administered consumers using the syntax in the following section.

MODIFY

```
modify consumer [guest-name]
parameter:
state [up|down|reset]
```

DISPLAY

```
list consumer
list consumer [guest-name]
```

DESCRIPTION

You can use the consumer component to manage the state of guest access to device health status services.

EXAMPLES

```
modify consumer tesseract-guest down
```

Transitions to 'down' the communication link between the hypervisor and guest pfmand daemons. When 'down', the guest is unable to send requests to the hypervisor, nor will it receive status updates.

```
modify consumer tesseract-guest up
```

Transitions to 'up' the communication link between the hypervisor and guest pfmand daemons. When 'up', the guest can send requests to the hypervisor, and receive status updates.

```
modify consumer tesseract-guest reset
```

Transitions to 'down', then to 'up', the communication link between the hypervisor and guest pfmand daemons.

```
list consumer tesseract-guest
```

Lists the current status of the guest called "tesseract-guest".

```
list consumer
```

Lists the current status of all guests.

SEE ALSO

list, modify, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013, 2015-2016. All rights reserved.

BIG-IP 2016-04-21 sys pfman consumer(1)

sys pfman device

NAME

device - Manage the state of a pfman controlled device status.

MODULE

sys pfman

SYNTAX

Manage pfman administered devices using the syntax in the following section.

MODIFY

```
modify device [pci_device]
parameter:
state [up|down|reset]
```

DISPLAY

```
list device
list device [pci_device]
```

DESCRIPTION

You can use the device component to manage the state of device reset and health status reporting.

EXAMPLES

```
modify device 04:00.0 down
```

This will cause pfmand to "down" the device. A device in the "down" state will not return to use until the device returns to "up" status. All associated tmm instances will be informed of the device status change to "down".

```
modify device 86.00.0 up
```

This causes pfmand to attempt to bring a "down" device back to the "up" status. When successful, all associated tmm instances will be informed of the device status change to "up".

```
modify device 86:00.0 reset
```

Instructs pfmand to issue a reset of device 86:00.0.

```
list device 85:00.0
```

Lists the status of device 85:00.0.

```
list device
```

Lists the status of all devices.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013, 2015-2016. All rights reserved.

BIG-IP 2016-04-21 sys pfman device(1)

sys proc-info

NAME

proc-info - Display CPU and memory usage for each process.

MODULE

sys

SYNTAX

Display proc-info component within the sys module using the syntax in the following section.

DISPLAY

```
show proc-info
show proc-info process_name
options:
(default | field-fmt | all | kil | meg | gig | raw | exa | peta | tera | zetta | yotta)
```

DESCRIPTION

Show proc-info displays CPU and memory usage for each process and the process associated module name. This can be used to debug which process or module uses more resource.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2011. All rights reserved.

BIG-IP 2011-02-16 sys proc-info(1)

sys provision

NAME

provision - Configures provisioning on the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the provision component within the sys module using the syntax in the following sections.

MODIFY

modify provision [afm | apm | asm | avr | cgnat | fps | gtm | ilx | lc | ltm | pem | swg | urldb | sslo | vcmp]

options:

cpu-ratio [integer]
disk-ratio [integer]
level [custom | dedicated | minimum | nominal | none]
memory-ratio [integer]

edit provision

[[[afm | apm | asm | avr | cgnat | fps | gtm | ilx | lc | ltm | pem | swg | urldb | sslo | vcmp] [glob] [regex]] ...]

options:

all-properties
non-default-properties

DISPLAY

list provision

[[[afm | apm | asm | avr | cgnat | fps | gtm | ilx | lc | ltm | pem | swg | urldb | sslo | vcmp] [glob] [regex]] ...]

show running-config provision

[[[afm | apm | asm | avr | cgnat | fps | gtm | ilx | lc | ltm | pem | swg | urldb | sslo | vcmp] [glob] [regex]] ...]

options:

all-properties
non-default-properties
one-line

DESCRIPTION

You can use the provision component to modify the allocation of resources to the licensed modules on your system.

EXAMPLES

modify provision asm level minimum

Provisions the minimum amount of resources for the BIG-IP Application Security Manager.

list provision

Displays the current provisioning of the system.

Using Transactions

1. create / cli transaction
2. modify / sys provision ltm level minimum
3. modify / sys provision gtm level nominal
4. submit / cli transaction

The previous four steps create a transaction to modify the provisioning of a unit to provision the Local Traffic Manager at the minimum level and the Global Traffic Manager at the nominal level.

1. create / cli transaction

2. modify / sys provision ltm level none
3. modify / sys provision gtm level dedicated
4. submit / cli transaction

The previous four steps create a transaction to modify the provisioning of a unit to dedicate all of the unit's resources to the Global Traffic Manager when currently only the Local Traffic Manager is provisioned.

OPTIONS

afm Specifies that you are provisioning the BIG-IP Advanced Firewall Manager. When the Advanced Firewall Manager is provisioned, the tmsh module security is enabled.

apm Specifies that you are provisioning the BIG-IP Access Policy Manager. When the Access Policy Manager is provisioned, the tmsh module apm is enabled.

asm Specifies that you are provisioning the BIG-IP Application Security Manager. When asm is provisioned the tmsh module asm is enabled.

avr Specifies that you are provisioning the BIG-IP Application Visibility and Reporting. When Application Visibility and Reporting is provisioned the tmsh module analytics is enabled.

cpu-ratio
Use this option only when the level option is set to custom. F5 Networks recommends that you do not modify this option. The default value is none.

disk-ratio
Use this option only when the level option is set to custom. F5 Networks recommends that you do not modify this option. The default value is none.

fps Specifies that you are provisioning the BIG-IP Fraud Protection Service. When fps is provisioned the tmsh module security anti-fraud is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

gtm Specifies that you are provisioning the BIG-IP Global Traffic Manager. When gtm is provisioned the tmsh module gtm is enabled.

ilx Specifies that you are provisioning BIG-IP iRules Language Extensions.

lc Specifies that you are provisioning the BIG-IP Link Controller. When Link Controller is provisioned the tmsh module gtm is enabled.

level
Specifies the level of resources that you want to provision for a module. The options are:

custom
F5 Networks does not recommend that you specify this level.

dedicated
Specifies that all resources are dedicated to the module you are provisioning. For all other modules, the level option must be set to none.

minimum
Specifies that you want to provision the minimum amount of resources for the module you are provisioning.

nominal
Specifies that you want to share all of the available resources equally among all of the modules that are licensed on the unit.

none Specifies that you do not want to provision any resources for this module.

ltm Specifies that you are provisioning the BIG-IP Local Traffic Manager.

memory-ratio
Use this option only when the level option is set to custom. F5 Networks recommends that you do not modify this option. The default value is none.

pem Specifies that you are provisioning the BIG-IP Policy Enforcement Manager. When Policy Enforcement Manager is provisioned the tmsh module pem is enabled.

regex
Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

swg Specifies that you are provisioning the BIG-IP Secure Web Gateway. When Secure Web Gateway is provisioned, the assumption is that Access Policy Management is already provisioned and tmsh component apm url-filter is enabled;

urldb
Specifies that you are provisioning the BIG-IP Secure Web Gateway with minimum resource mode. Mainly used for SSL orchestration use cases. The tmsh module urldb cannot be provisioned in BIG-IP when the tmsh module swg is already provisioned and vice-versa.

sslo Specifies that you are provisioning the BIG-IP SSL Otrchestrator.

vcmp Specifies that you are provisioning the BIG-IP Virtual CMP. When Virtual CMP is provisioned the tmsh module vcmp is enabled.

SEE ALSO

edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2020-02-12 sys provision(1)

sys pva-traffic

NAME

pva-traffic - Displays and resets Packet Velocity(r) ASIC (PVA) traffic statistics for the system.

MODULE

sys

SYNTAX

Configure the pva-traffic component within the sys module using the following syntax.

MODIFY

reset-stats pva-traffic

DISPLAY

show pva-traffic

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

global

DESCRIPTION

You can use the pva-traffic component to display traffic statistics, including bits in and out, packets in and out, current, maximum, and total connections, and other miscellaneous statistics.

The BIG-IP(r) system has one PVA accelerator; however, when you run the command show pva-traffic, the system displays a PVA statistics entry for each Traffic Management Microkernel (TMM).

EXAMPLES

show pva-traffic

Displays PVA traffic statistics for the system.

show pva-traffic raw

Displays PVA traffic statistics for the system in raw data form.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, sys tmm-traffic, sys traffic, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 sys pva-traffic(1)

sys raid array

NAME

array - Configures an array of hard disks on the BIG-IP(r) system.

MODULE

sys raid

SYNTAX

Configure the array component within the sys raid module using the syntax in the following sections.

MODIFY

modify array [name] [[add | remove] [hard disk name]]

DISPLAY

show array
show array [name]
options:
field-fmt

DESCRIPTION

You can use the array component to add a hard disk to or remove a hard disk from an array of disks, or to display information about an array of disks.

EXAMPLES

show array

Displays information about all of the arrays that are configured on the system.

modify array MD1 remove HD2

Removes hard disk, HD2 from array, MD1.

OPTIONS

hard disk name

Specifies the name of the hard disk that you want to add to or remove from the array. This option is required for the command modify.

name Specifies the name of the array. This option is required for the command modify.

SEE ALSO

modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-11 sys raid array(1)

sys raid bay

NAME

bay - Manages a BIG-IP(r) system disk drive bay.

MODULE

sys raid

SYNTAX

Manage the bay component within the sys raid module using the syntax in the following sections.

MODIFY

modify bay [1 | 2]
options:
flash-led
no-flash-led

DISPLAY

show bay [1 | 2]
options:
field-fmt

DESCRIPTION

You can use the bay component to display information about a system bay, signal the LED on a bay to flash, or signal the LED on a bay to stop flashing. The LED is helpful for identifying the location of a specific disk, see sys raid disk.

EXAMPLES

modify bay 1 flash-led

Signal the system to make the LED on bay 1 flash.

show bay

Displays information about the system bay.

show bay field-fmt

Displays information about the system bay in a field format.

OPTIONS

flash-led

Signal the LED on the bay to flash.

no-flash-led

Signal the LED on the bay to stop flashing.

For information about the field-fmt option, see help show.

SEE ALSO

show, sys raid disk, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2011-08-29 sys raid bay(1)

sys raid disk

NAME

disk - Displays information about the BIG-IP(r) system disks.

MODULE

sys raid

SYNTAX

Display information about the disk component within the sys raid module using the syntax in the following sections.

DISPLAY

show disk [name]

options:

field-fmt

all-properties

DESCRIPTION

You can use the disk component to display information about the system disks including name, serial number, and whether the disk is a member of an array of disks. When "all-properties" option is specified, the media wear-out information of the disk is also shown. This include the wear-out indicator, space available, power-on hours, and estimated remaining life.

EXAMPLES

show disk

Displays information about all of the system disks.

show disk HD1 field-fmt

Displays information, in a field format, about disk, HD1.

show disk SSD1 all-properties

Displays all information (including the media wear-out information) about disk, SSD1.

OPTIONS

name Specifies the name of the disk for which you want to display information.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

sys ready

NAME

ready - Display the ready status for processing config/license/provision commands.

MODULE

sys

SYNTAX

Display the ready status for processing config/license/provision commands using the syntax in the following section.

DISPLAY

show ready
options:
field-fmt

DESCRIPTION

You can use the ready component to display whether or not the BIG-IP is ready to handle actions such as:

config

Adding or modifying the system's configuration, such as adding a pool member.

license

Apply or update licensing.

provision

Modify resource provisioning, such as enabling or disabling modules

EXAMPLES

show ready

```
-----  
Sys::BIG-IP Ready:  
-----
```

```
config  yes  
license yes  
provision yes
```

```
show ready field-fmt
```

```
sys bigip-ready {  
  config-ready yes  
  license-ready yes  
  provision-ready yes  
}
```

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

sys scripted

NAME

scriptd - Configure the scriptd daemon

MODULE

sys

SYNTAX

Configure the scriptd daemon within the sys module using the syntax in the following sections.

MODIFY

modify scriptd

options:

log-level [alert | crit | debug | emerg | err | info | notice | warn]

max-script-run-time [seconds]

DISPLAY

list scriptd

show running-config scriptd

options:

all-properties

DESCRIPTION

You can use the scriptd component to configure the scriptd daemon. The scriptd daemon runs app application template implementation scripts when an application service is created or updated (see sys application template and sys application service).

EXAMPLES

list scriptd

Displays scriptd configuration.

modify scriptd max-script-run-time 120

Updates the maximum time, in seconds, that a script is allowed to run.

OPTIONS

log-level

Specifies the syslog level at which scriptd will generate log messages.

max-script-run-time

Specifies, in seconds, the maximum amount of time that a script is allowed to run before scriptd will kill the script. The default value is 300. The minimum value is 5.

SEE ALSO

list, modify, show, sys application template, sys application service, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP **2012-04-19** **sys scriptd(1)**

sys service

NAME

service - Manages services on the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the service component within the sys module using the syntax in the following sections.

MODIFY

modify service [name]

options:

[add | disable | enable | reinit | remove]

restart service [name]

start service [name]

stop service [name]

options:

force

DISPLAY

list service

list service [name]

show running-config service

show running-config service [name]

options:

all-properties

show service

options:
memstat

DESCRIPTION

You can use the service component to add, disable or enable, start, stop, restart, reinitialize, remove, or display information about a service.

Note that the tmsh connection to mcpd will be dropped if you stop or restart the mcpd service. The next tmsh command will prompt you to try again. Alternatively you can quit tmsh and login again.

EXAMPLES

list service

Displays information about the services available on the BIG-IP system.

restart service mcpd

Restarts the mcpd daemon.

OPTIONS

add Adds the specified service.

disable

Disables the specified service.

enable

Enables the specified service.

memstat

Displays memory usage statistics for the specified service.

reinit

Reinitializes the specified service.

remove

Removes the specified service.

SEE ALSO

list, modify, restart, show, start, stop, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2011-04-29 sys service(1)

sys sflow data-source http

NAME

http - Displays the status of all HTTP sFlow data sources on the BIG-IP system.

MODULE

sys sflow data-source

SYNTAX

Display the status of http component within the sys sflow data-source module using the syntax shown in the following sections.

DISPLAY

show http
options:
all-properties
field-fmt

DESCRIPTION

You can use the http component to display the current status of all HTTP sFlow data sources on the BIG-IP system.

EXAMPLES

show http

Displays the current status of all HTTP sFlow data sources.

SEE ALSO
show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-12-28 sys sflow data-source http(1)

sys sflow data-source interface

NAME
interface - Displays the status of all sFlow data sources (interfaces) on the BIG-IP system.

MODULE
sys sflow data-source

SYNTAX
Display the status of interface component within the sys sflow data-source module using the syntax shown in the following sections.

DISPLAY
show interface
options:
all-properties
field-fmt

DESCRIPTION
You can use the interface component to display the current status of all sFlow data sources (interfaces) on the BIG-IP system.

EXAMPLES
show interface

Displays the current status of all sFlow data sources (interfaces).

SEE ALSO
show, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-12-28 sys sflow data-source interface(1)

sys sflow data-source system

NAME
system - Displays the status of the system sFlow data sources on the BIG-IP system.

MODULE
sys sflow data-source

SYNTAX
Display the status of system component within the sys sflow data-source module using the syntax shown in the following sections.

DISPLAY
show system
options:
all-properties
field-fmt

DESCRIPTION
You can use the system component to display the current status of the system sFlow data sources on the BIG-IP system.

EXAMPLES

show system

Displays the current status of the system sFlow data sources.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-12-28 sys sflow data-source system(1)

sys sflow data-source vlan

NAME

vlan - Displays the status of all sFlow data sources (VLANs) on the BIG-IP system.

MODULE

sys sflow data-source

SYNTAX

Display the status of vlan component within the sys sflow data-source module using the syntax shown in the following sections.

DISPLAY

show vlan
options:
all-properties
field-fmt

DESCRIPTION

You can use the vlan component to display the current status of all sFlow data sources (VLANs) on the BIG-IP system.

EXAMPLES

show vlan

Displays the current status of all sFlow data sources (VLANs).

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2012. All rights reserved.

BIG-IP 2012-12-28 sys sflow data-source vlan(1)

sys sflow global-settings http

NAME

http - Manages the global HTTP sFlow configuration on the BIG-IP system.

MODULE

sys sflow global-settings

SYNTAX

Configure the http component within the sys sflow global-settings module using the syntax shown in the following sections.

MODIFY

modify http
options:

description [string]
poll-interval [integer]
sampling-rate [integer]

DISPLAY

list http
options:
all-properties
non-default-properties
one-line

DESCRIPTION

You can use the http component to modify or list the global HTTP sFlow configuration on the BIG-IP system.

Note: You can modify the global HTTP sFlow configuration on the BIG-IP system, only if you are assigned either the Resource Administrator or Administrator user role.

EXAMPLES

modify http poll-interval 60 sampling-rate 1500

Sets the poll-interval to 60 seconds and the sampling-rate to 1500 packets for all monitored HTTP data sources on the BIG-IP system.

OPTIONS

description
User defined description.

poll-interval
Specifies the maximum interval in seconds between polling by the sFlow agent of all monitored HTTP data sources on the BIG-IP system. The default value is 10.

sampling-rate
Specifies the ratio of packets observed at all HTTP data sources to the samples generated. For example, a sampling rate of 2000 specifies that 1 sample will be randomly generated for every 2000 packets observed. The default value is 1024.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-10-15 sys sflow global-settings http(1)

sys sflow global-settings interface

NAME

interface - Manages the global sFlow configuration for interfaces on the BIG-IP system.

MODULE

sys sflow global-settings

SYNTAX

Configure the interface component within the sys sflow global-settings module using the syntax shown in the following sections.

MODIFY

modify interface
options:
description [string]
poll-interval [integer]

DISPLAY

list interface
options:
all-properties
non-default-properties
one-line

DESCRIPTION

You can use the interface component to modify or list the global sFlow configuration for interfaces on the BIG-IP system.

Note: You can modify the global sFlow configuration for interfaces on the BIG-IP system, only if you are assigned either the Resource Administrator or Administrator user role.

EXAMPLES

```
modify interface poll-interval 60
```

Sets the poll-interval to 60 seconds for all monitored data sources (interfaces) on the BIG-IP system.

OPTIONS

description
User defined description.

poll-interval
Specifies the maximum interval in seconds between polling by the sFlow agent of all monitored data sources (interfaces) on the BIG-IP system. The default value is 10.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2013-10-15 sys sflow global-settings interface(1)

sys sflow global-settings system

NAME

system - Manages the global system sFlow configuration on the BIG-IP system.

MODULE

sys sflow global-settings

SYNTAX

Configure the system component within the sys sflow global-settings module using the syntax shown in the following sections.

MODIFY

```
modify system
options:
  description [string]
  poll-interval [integer]
```

DISPLAY

```
list system
options:
  all-properties
  non-default-properties
  one-line
```

DESCRIPTION

You can use the system component to modify or list the global system sFlow configuration on the BIG-IP system.

Note: You can modify the global system sFlow configuration on the BIG-IP system, only if you are assigned either the Resource Administrator or Administrator user role.

EXAMPLES

```
modify system poll-interval 60
```

Sets the poll-interval to 60 seconds for the system data sources on the BIG-IP system.

OPTIONS

description
User defined description.

poll-interval
Specifies the maximum interval in seconds between polling by the sFlow agent of the system data sources on the BIG-IP system. The default value is 10.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

sys sflow global-settings vlan

NAME

vlan - Manages the global sFlow configuration for VLANs on the BIG-IP system.

MODULE

sys sflow global-settings

SYNTAX

Configure the vlan component within the sys sflow global-settings module using the syntax shown in the following sections.

MODIFY

modify vlan

options:

description [string]

poll-interval [integer]

sampling-rate [integer]

DISPLAY

list vlan

options:

all-properties

non-default-properties

one-line

DESCRIPTION

You can use the vlan component to modify or list the global sFlow configuration for VLANs on the BIG-IP system.

Note: You can modify the global sFlow configuration for VLANs on the BIG-IP system, only if you are assigned either the Resource Administrator or Administrator user role.

EXAMPLES

modify vlan poll-interval 60 sampling-rate 1500

Sets the poll-interval to 60 seconds and the sampling-rate to 1500 packets for all monitored data sources (VLANs) on the BIG-IP system.

OPTIONS

description

User defined description.

poll-interval

Specifies the maximum interval in seconds between polling by the sFlow agent of all monitored data sources (VLANs) on the BIG-IP system. The default value is 10.

sampling-rate

Specifies the ratio of packets observed at all data sources (VLANs) to the samples generated. For example, a sampling rate of 2000 specifies that 1 sample will be randomly generated for every 2000 packets observed. The default value is 2048.

SEE ALSO

list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

sys sflow receiver

NAME

receiver - Manages sFlow receivers configured on the BIG-IP system.

MODULE

sys sflow

SYNTAX

Configure the receiver component within the sys sflow module using the syntax shown in the following sections.

CREATE/MODIFY

```
create receiver [name]
modify receiver [name]
options:
  address [ip address]
  app-service [[string] | none]
  description [string]
  max-datagram-size [integer]
  port [ip port]
  state [disabled | enabled]
```

DISPLAY

```
list receiver
list receiver [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  one-line
```

DELETE

```
delete receiver [name]
```

DESCRIPTION

You can use the receiver component to create, delete, list, or modify an sFlow receiver object on the BIG-IP system.

Note: You can add an sFlow receiver to the BIG-IP system, only if you are assigned either the Resource Administrator or Administrator user role.

EXAMPLES

```
create receiver my_receiver address 10.10.10.10
```

Creates an sFlow receiver object named `my_receiver` with an IP address of `10.10.10.10`, where the port, max-datagram-size, and state options are set to default values.

```
create receiver my_receiver address 10.20.10.20 port 1234 state enabled
```

Creates an sFlow receiver object named `my_receiver` with an IP address of `10.20.10.20`, a port of `1234`, and the max-datagram-size option set to default value. The state of the receiver is enabled.

```
modify receiver my_receiver state enabled
```

Changes the state of sFlow receiver object named `my_receiver` to enabled.

OPTIONS

address

Specifies the IP address on which the sFlow receiver listens for UDP datagrams. This option is required for the create command.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

max-datagram-size

Specifies the maximum size in bytes of the UDP datagram the sFlow receiver accepts. The default value is 1400.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

port Specifies the port on which the sFlow receiver listens for UDP datagrams. The default value is the standard sFlow port, 6343.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@`[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

state

Specifies the state of the receiver. The sFlow samples will be collected and sent to the receiver when enabled. The default value is disabled.

SEE ALSO

create, delete, glob, list, modify, regex, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2014-03-18 sys sflow receiver(1)

sys smtp-server

NAME

smtp-server - Configure the SMTP server connection.

MODULE

sys

SYNTAX

Create or modify an SMTP server access configuration using the syntax in the following sections.

CREATE / MODIFY

modify smtp-server [name]

create smtp-server [name]

options:

app-service [[string] | none]
[authentication-enabled | authentication-disabled]
encrypted-connection [none | tls | ssl]
local-host-name [string]
smtp-server-host-name [string]
smtp-server-port [integer]
from-address [string]
username [string]
password [string]

DISPLAY

list smtp-server

show running-config smtp-server

options:

all-properties

DESCRIPTION

You can use the smtp-server component to configure an SMTP server connection.

EXAMPLES

list smtp-server

Displays the SMTP configuration.

modify smtp-server smtp1 authentication-enabled encrypted-connection ssl local-host-name example.f5.com from-address example@f5.com smtp-server-host-name mail.server.com username user password pass

Configures SMTP server connection with username=user and password=pass to be authenticated against the SMTP server mail.server.com. SSL encryption will be used for all communication with the SMTP server. Email messages will be sent out with the address example@f5.com in the "Reply-To" address.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

[authentication-enabled | authentication-disabled]

Enables or disables authentication against the configured SMTP server.

encrypted-connection

Specifies which type of encrypted connection the SMTP server requires in order to send mail. The default value is none.

local-host-name

Specifies the host name used in SMTP headers in the format of a fully qualified domain name. This setting does not refer to the BIG-IP system's Hostname.

smtp-server-host-name

Specifies the SMTP server host name in the format of a fully qualified domain name.

smtp-server-port

Specifies the SMTP port number. The default value is 25.

from-address

Specifies the email address that the email is being sent from. This is the "Reply-to" address that the recipient sees.

username
Specifies the user name that the SMTP server requires when validating a user.

password
Specifies the password that the SMTP server requires when validating a user. This password is stored in an encrypted form.

SEE ALSO

list, create, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-05-22 sys smtp-server(1)

sys snmp

NAME

snmp - Configures the simple network management protocol (SNMP) daemon for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the snmp component within the sys module using the following syntax.

MODIFY

modify snmp

options:

```
agent-addresses [add | delete | replace-all-with] {
    ["agent:port"] ...
}
agent-addresses none
agent-trap [enabled | disabled]
allowed-addresses [add | delete | replace-all-with] {
    [IP address]
}
allowed-addresses none
auth-trap [enabled | disabled]
bigip-traps [enabled | disabled]
communities [add | delete | modify | replace-all-with] {
    [name] {
        options:
access [ro | rw]
community-name [string]
description [string]
ipv6 [enabled | disabled]
oid-subset [string]
source [ [ip address] | [FQDN] | [ [protocol]:[ip address] ] |
    [ [protocol]:[FQDN] ] ]
    }
}
communities none
description [string]
disk-monitors [add | delete | modify | replace-all-with] {
    [name] {
        options:
description [string]
minspace [integer]
minspace-type [percent | size]
path [string]
    }
}
disk-monitors none
include [string]
l2forward-vlan [all | add | delete | replace-all-with] {
    [VLAN name] ...
}
l2forward-vlan none
load-max1 [integer]
load-max5 [integer]
load-max15 [integer]
```

```

process-monitors [add | delete | modify | replace-all-with] {
  [name] {
    options:
description [string]
process [string]
min-processes [integer]
max-processes [ [integer] | infinity ]
  }
}
process-monitors none
snmpv1 [enabled | disabled]
snmpv2 [enabled | disabled]
sys-contact [string]
sys-location [string]
sys-services [integer]
trap-community [string]
trap-source [IP address]
traps [add | delete | modify | replace-all-with] {
  [name] {
    options:
auth-password [string]
auth-protocol [md5 | sha | none]
community [string]
description [string]
engine-id [ [number] | none ]
host [ [ip address] | [FQDN] | [ [protocol]:[ip address] ] |
  [ [protocol]:[FQDN] ] ]
port [integer]
privacy-password [string]
privacy-protocol [aes | des | none]
security-level [auth-no-privacy | auth-privacy | no-auth-no-privacy]
security-name [string]
version [1 | 2c | 3]
  }
}
traps none
users [add | delete | modify | replace-all-with] {
  [user name] {
    options:
access [ro | rw]
auth-password [string]
auth-protocol [md5 | sha | none]
description [string]
oid-subset [string]
privacy-password [string]
privacy-protocol [aes | des | none]
security-level [auth-no-privacy | auth-privacy | no-auth-no-privacy]
username [string]
  }
}
v1-traps [add | delete | modify | replace-all-with] {
  [name] {
    options:
community [string]
description [string]
host [ [ip address] | [FQDN] | [ [protocol]:[ip address] ] |
  [ [protocol]:[FQDN] ] ]
port [integer]
  }
}
v1-traps none
v2-traps [add | delete | modify | replace-all-with] {
  [name] {
    options:
community [string]
description [string]
host [ [ip address] | [FQDN] | [ [protocol]:[ip address] ] |
  [ [protocol]:[FQDN] ] ]
port [integer]
  }
}
v2-traps none

edit snmp
options:
  all-properties
  non-default-properties

DISPLAY
list snmp
list snmp [option]
show running-config snmp
show running-config snmp [option]
options:
  all-properties
  non-default-properties
  one-line

```

DESCRIPTION

You can use the `snmp` component to configure the `snmpd` daemon for the BIG-IP system.

Important: F5 Networks recommends that users of the Configuration utility exit the utility before changes are made to the system using the command sequence `tmsch sys snmp`. This is because making changes to the system using this command causes a restart of the `snmpd` daemon. Likewise, restarting the `snmpd` daemon creates the necessity for a restart of the Configuration utility.

EXAMPLES

```
modify snmp sys-contact admin@company.com
```

Modifies the configuration to indicate that the person who administers the `snmpd` daemon for the system can be reached using the email address, `admin@company.com`.

```
modify snmp sys-location "central office"
```

Modifies the configuration to indicate that the physical location of the system is the central office.

```
modify snmp snmpv1 disabled
```

Disables `snmpV1` agent support.

```
modify snmp snmpv2c disabled
```

Disables `snmpV2c` agent support.

```
modify snmp agent-trap disabled
```

Disables agent traps.

```
modify snmp allowed-addresses add {10.10.0.0/255.255.240.0}
```

Adds a range of SNMP clients to the `/etc/hosts.allow` file.

```
modify snmp traps add { tv1 { version 1 community public host 192.168.1.240 port 162 } }
```

Adds an SNMP version 1 trapsess, `tv1`, to the system. The destination IP address of `tv1` is `192.168.1.240`, the port is `162`, and the community that has access to `tv1` is `public`. The default port is `162`.

```
modify snmp traps add { tv2 {version 2c community public host 192.168.1.241 port 162} }
```

Adds an SNMP version 2 trapsess, `tv2`, to the system. The destination IP address of `tv2` is `192.168.1.241`, the port is `162`, and the community that has access to `tv2` is `public`. The default port is `162`. The default version is `2c` (version 2).

```
modify snmp traps add { trap_v3_1 { version 3 host 192.168.1.242 port 162 security-level auth-no-privacy security-name mySecurityName auth-protocol md5 auth-password myAuthPassword } }
```

Adds an SNMP version 3 trapsess, `trap_v3_1`, with authentication capabilities to the system. The destination IP address of `trap_v3_1` is `192.168.1.242`, the port is `162`, the security level is the authentication without privacy, the security name is `mySecurityName`, the authentication protocol is `MD5`, and the authentication password is `myAuthPassword`. The default port is `162`.

```
modify snmp traps add { trap_v3_2 { version 3 host 192.168.1.243 port 162 security-level auth-privacy security-name mySecurityName auth-protocol sha auth-password myAuthPassword privacy-protocol aes privacy-password myPrivacyPassword } }
```

Adds an SNMP version 3 trapsess, `trap_v3_2`, with authentication and privacy capabilities to the system. The destination IP address of `trap_v3_2` is `192.168.1.243`, the port is `162`, the security level is the authentication and privacy, the security name is `mySecurityName`, the authentication protocol is `SHA`, the authentication password is `myAuthPassword`, the privacy protocol is `AES`, and the privacy password is `myPrivacyPassword`. The default port is `162`.

```
modify snmp v1-traps add { ts { community public host 10.20.5.11 port 162 } }
```

Adds an SNMP version 1 trapsink, `ts`, to the system. The destination IP address of `ts` is `10.20.5.11`, the port is `162`, and the community that has access to `ts` is `public`. The default port is `162`.

```
modify snmp v2-traps add { t2s { community public host 10.20.5.12 port 162 } }
```

Adds an SNMP version 2 trap2sink, `t2s`, to the system. The destination IP address of `t2s` is `10.20.5.12`, the port is `162`, and the community that has access to `t2s` is `public`. The default port is `162`.

```
modify snmp users add { myUser1 { username myUser1 access ro security-level auth-no-privacy auth-protocol md5 auth-password myAuthPassword privacy-protocol } }
```

Adds an SNMP version 3 user with the user name, `myUser1`, to the system. The access to the management information base (MIB) of `myUser1` is read-only, the security level is the authentication without privacy, the authentication protocol is `MD5`, and the authentication password is `myAuthPassword`.

```
modify snmp users add { myUser2 { username myUser2 oid-subset .1.3.6.1.4.1.3375 auth-protocol md5 auth-password myAuthPassword privacy-protocol none } }
```

Adds an SNMP version 3 user with the user name, `myUser2`, to the system. The access to the management information base (MIB) of `myUser2` is read-only (by default) and restricted to every object below `.1.3.6.1.4.1.3375` object identifier in the MIB tree, the security level is the authentication without privacy,

the authentication protocol is MD5, and the authentication password is myAuthPassword.

```
modify snmp users add { myUser3 { username myUser3 access ro security-level auth-privacy auth-protocol sha  
auth-password myAuthPassword privacy-protocol des privacy-password myPrivacyPassword } }
```

Adds an SNMP version 3 user with the user name, myUser3, to the system. The access to the management information base (MIB) of myUser3 is read-only, the security level is the authentication and privacy, the authentication protocol is SHA, the authentication password is myAuthPassword, the privacy protocol is DES, and the privacy password is myPrivacyPassword.

```
modify snmp users add { myUser4 { username myUser4 access ro security-level no-auth-no-privacy auth-protocol  
none privacy-protocol none } }
```

Adds an SNMP version 3 user with the user name, myUser4, to the system. The access to the management information base (MIB) of myUser4 is read-only without the authentication and privacy settings.

```
modify snmp communities add { community1 { community-name mycommunity access ro source 192.168.1.246 oid-  
subset 5 ipv6 disabled } }
```

Creates a community specification named community1 for the BIG-IP system. community1 includes a community, named mycommunity, that provides read-only access to the host at 192.168.1.246. This host cannot be an IPv6 address. The oid for this community is 5.

```
modify snmp communities add { new-name { community-name public source default oid-subset 1 access ro } }
```

Replaces the default community specification for the BIG-IP system. Using this command, the default community includes a community, named public, that provides read-only access to the default host. The oid for this community is 1.

```
modify snmp communities delete { mycommunity }
```

Deletes the community named mycommunity.

```
modify snmp load-max1 0 load-max5 0 load-max15 0
```

Disables monitoring of snmpd load average on the BIG-IP system.

OPTIONS

snmpv1

Specifies, when enabled, that the snmpd daemon supports snmpV1 queries. The default value is enabled.

snmpv2c

Specifies, when enabled, that the snmpd daemon supports snmpV2c queries. The default value is enabled.

agent-addresses

Indicates that the SNMP agent is to listen on the specified address. F5 Networks recommends that you do not change this setting without fully understanding the impact of the change.

agent-trap

Specifies, when enabled, that the snmpd daemon sends traps, for example, start and stop traps. The default value is enabled.

allowed-addresses

Configures the IP addresses of the SNMP clients from which the snmpd daemon accepts requests. An SNMP client is a system that runs the SNMP manager software for the purpose of remotely managing the BIG-IP system. The default value is 127.

auth-trap

Specifies, when enabled, that the snmpd daemon generates authentication failure traps. The default value is disabled.

bigip-traps

Specifies, when enabled, that the BIG-IP system sends device warning traps to the trap destinations. The default value is enabled.

community

Configures a community for the snmpd daemon. Note that you must include a community key, and you must enclose the attributes in braces.

The options are additive and include:

access

Specifies the community access level to the MIB. The access options are ro (read-only) or rw (read-write). The default value is ro.

community name

Specifies the name of the community that you are configuring for the snmpd daemon. This option is required. The default value is public.

description

User defined description.

ipv6 Specifies to enable or disable IPv6 addresses for the community that you are configuring. The default value is disabled.

oid-subset

Specifies to restrict access by the community to every object below the specified object identifier

(OID).

source

Specifies the source addresses with the specified community name that can access the management information base (MIB). The default value is default, which means allow any source address to access the MIB.

description

User defined description.

disk-monitors

Checks the disks mounted at the specified path for available disk space.

The options are:

description

User defined description.

minspace

Specifies the minimum disk space threshold in either kB or percentage based on the value of the minspace-type option. If the available disk space is less than this amount, the associated entry in the 1.3.6.1.4.1.2021.9.1.100 MIB table is set to (1) and a descriptive error message is returned to queries of 1.3.6.1.4.1.2021.9.1.101.

minspace-type

Specifies a minimum disk space measurement type of either size in kB, or percent. Note that the value of the minspace option is based on the value of this option.

path Specifies the path to the disk that the system checks for disk space. This option is required.

include

Warning: Do not use this parameter without assistance from the F5 Technical Support team. The system does not validate the commands issued using the include parameter. If you use this parameter incorrectly, you put the functionality of the system at risk.

l2forward-vlan

Specifies the VLANs for which you want the snmpd daemon to expose Layer 2 forwarding information. Layer 2 forwarding is the means by which frames are exchanged directly between hosts, with no IP routing required. The default value is none.

The options are:

all The snmpd daemon exposes Layer 2 forwarding information for all VLANs.

Warning: When you set this option to all, the system can create a very large table of statistics and potentially affect system performance.

none Indicates that this option is not set.

Important: The default is not the same as setting this option to the string "none," which indicates that you do not want the snmpd daemon to expose Layer 2 forwarding for any VLAN.

VLAN name

Specifies the names of the VLANs for which the snmpd daemon exposes Layer 2 forwarding information. The snmpd daemon overwrites the value of the sysL2ForwardAttrVlan object identifier (OID) with the specified VLAN names. Once you set this parameter, users cannot change the value of the sysL2ForwardAttrVlan OID using the SNMP set method.

load-max1

Specifies the maximum 1-minute load average of the machine. If the load exceeds this threshold, the associated entry in the 1.3.6.1.4.1.2021.10.1.100 MIB table is set to (1) and a descriptive error message is returned to queries of 1.3.6.1.4.1.2021.10.1.101.

Note that when you specify a 0 (zero) for all three of the load-max1, load-max5, and load-max15 options, the system does not monitor the load average.

load-max5

Specifies the maximum 5-minute load average of the machine. If the load exceeds this threshold, the associated entry in the 1.3.6.1.4.1.2021.10.1.100 MIB table is set to (1) and a descriptive error message is returned to queries of 1.3.6.1.4.1.2021.10.1.101.

Note that when you specify a 0 (zero) for all three of the load-max1, load-max5, and load-max15 options, the system does not monitor the load average.

load-max15

Specifies the maximum 15-minute load average of the machine. If the load exceeds this threshold, the associated entry in the 1.3.6.1.4.1.2021.10.1.100 MIB table is set to (1) and a descriptive error message is returned to queries of 1.3.6.1.4.1.2021.10.1.101.

Note that when you specify a 0 (zero) for all three of the load-max1, load-max5, and load-max15 options, the system does not monitor the load average.

process-monitors

Specifies to check the machine to determine if the specified process is running. An error flag (1) and a description message are passed to the 1.3.6.1.4.1.2021.2.1.100 and 1.3.6.1.4.1.2021.2.1.101 MIB columns (respectively) if the specified program is not found in the process table as reported by /bin/ps -e.

F5 Networks recommends that you do not modify or delete system processes; however, you can add, modify, or delete user-defined processes.

The options are:

description

User defined description.

max-processes

Specifies the maximum number of instances of the process that can run. The default value is 1.

If you do not specify values for the min-processes and max-processes options, the max-processes option is 1 by default.

min-processes

Specifies the minimum number of instances of the process that can run. The default value is 1.

If you do not specify a value for the max-processes option, and the min-processes option is not specified, the min-processes option is 0 (zero) by default.

process

Specifies the name of the monitored process. The maximum length for a process name is 16 characters. This option is required.

sys-contact

Specifies the name of the person who administers the snmpd daemon for this system. The default value is "Customer Name"

sys software block-device-hotfix

NAME

block-device-software-hotfix - Manages F5 Networks block device software hotfixes.

MODULE

sys software

SYNTAX

Install or display information about a block-device-hotfix using the syntax in the following sections.

INSTALL

install block-device-hotfix [name] volume [name]

options:

create-volume
reboot

DISPLAY

list block-device-hotfix

list block-device-hotfix [[[name [/slot_id]] | [glob] | [regex]] ...]

options:

build
checksum
device-agent
id
one-line
product
title
resource-id
verified
version

DESCRIPTION

You can use the block-device-hotfix component to install a block-device-hotfix onto a volume, or view information about available block-device-hotfixes.

Use the create-volume option with the block-device-hotfix component to create new volumes.

Note: You use the slot_id option only for chassis systems and only when displaying the values for the options of a specific block-device-hotfix. You do not use the slot_id option when installing or deleting a block-device-hotfix, because these commands operate on all blades or the entire system.

EXAMPLES

list block-device-hotfix Hotfix-BIGIP-11.4.0-2419.0-HF3.iso

Displays information about the specified block-device-hotfix, BIGIP-9.6.1-824.0-HF3.im.

list block-device-hotfix */1

Displays information about the all the block-device-hotfixes on the first slot.

install block-device-hotfix Hotfix-BIGIP-11.4.0-2419.0-HF3.iso volume HD1.1 reboot

Attempts to install the specified block-device-hotfix, Hotfix-BIGIP-11.4.0-2419.0-HF3.iso, onto HD1.1.

Note: If the installation is successful, and you used the reboot option, as in this example, the machine reboots into the newly installed block-device-hotfix.

OPTIONS

build

Displays the build number of the block-device-hotfix.

checksum

Displays the checksum of the block-device-hotfix.

create-volume

Create a new volume using the name specified with the volume option. Mirrored volume names must begin with the prefix MD1.. Mirrored volumes are available only on systems that support RAID, see sys raid.

device-agent

Displays the name of the service which is responsible for managing the type of block device on which a hotfix is available. vcmp-virtual-cdrom is the device-agent which allows VCMF guests to install hotfixes that reside in the host system, via a virtual cdrom device.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies the name and sequential ID of the block-device-hotfix that you want to install or delete.

product

Displays the F5 Networks product this block-device-hotfix contains.

reboot

Specifies that the system reboots immediately after a successful installation.

resource-id

Displays the resource-id string corresponding to software block-device-hotfix in question. This string is used to identify the image for use in interacting with the device agent that is responsible for management of the resource. In the case of the vcmp-virtual-cdrom device-agent the resource-id is used by the live installation daemon when requesting that the hypervisor make hotfixes available before the installation can proceed.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

slot id

Specifies the number of the slot on a chassis system that contains the block-device-hotfix about which you want to display information.

title

Displays a textual description of the block-device-hotfix.

verified

When set to yes, indicates that the block-device-hotfix is authentic.

version

Displays the version number of the product the block-device-hotfix contains.

volume

Specifies the name of the volume on which you want to install the block-device-hotfix, or from which you want to delete the block-device-hotfix.

SEE ALSO

delete, glob, install, list, regex, sys software block-device-image, sys software hotfix, sys software image, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

sys software block-device-image

NAME

block-device-image - Manages F5 Networks block device software images.

MODULE

sys software

SYNTAX

Install or display information about a block-device-image using the syntax in the following sections.

INSTALL

install

options:

create-volume
image [name]
reboot
volume [name]

DISPLAY

list block-device-image

list block-device-image [[[name [/slot_id]] | [glob] | [regex]] ...]

options:

build
build-date
checksum
device-agent
file-size
last-modified
one-line
product
resource-id
verified
version

DESCRIPTION

You can use the block-device-image component to install images from block devices onto a volume, or view information about available block-device-images.

INSTALLING A SOFTWARE BLOCK-DEVICE-IMAGE

Before you begin installing a block-device-image, the image must be made available to the system. As of this writing, block device images are only available from within a VCMP guest via a virtual cdrom service. A VCMP host administrator must have provided the images for use by the guest as part of their administration of the VCMP host.

From tmsh, you can use show sys software status to see all of the available disk volumes where you can install the image. You can install the image file in any volume that is not active.

Then use the install command with this component to install the image to an unused volume. You can use the create-volume option if you want to create a new volume. The installation takes some time; you can use show sys software status repetitively to watch the progress of the installation. To put the newly installed software into active service, use the reboot option in the install command, or use the reboot volume vol-name command after the install command completes.

Note: You use the slot_id option only for chassis systems and only when displaying the values for the options of a specific block-device-image. You do not use the slot_id option when installing or deleting a block-device-image, because these commands operate on all blades or the entire system.

CONFIRMING AN BLOCK-DEVICE-IMAGE INSTALLATION

You can use show sys version to confirm that the system is running the new software version. If this is a new module for the current system, you may need to use show sys license and/or install sys license to update your license. For a new module, you may also need to provision CPU, memory, and disk space for the module with the sys provision component.

EXAMPLES

```
install block-device-image BIGIP-11.4.1.608.0.iso volume HD1.1 reboot
```

Attempts to install the specified block-device-image, BIGIP-11.4.1.608.0.iso, onto HD1.1. **Note:** If the installation is successful, the machine reboots into the newly installed block-device-image.

```
list block-device-image BIGIP-11.4.1.608.0.iso
```

Displays information about the specified block-device-image, build 608.0 of BIG-IP version 11.4.1.

```
list block-device-image */1
```

Displays information about all of the block-device-images located on the first slot.

OPTIONS

build

Displays the build number of the block-device-image.

build-date

Displays the date on which the block-device-image was built.

checksum

Displays the checksum of the block-device-image. You can use this option to verify the integrity of the block-device-image.

create-volume

Creates a new volume using the name specified with the volume option. Mirrored volume names must begin with the prefix MD1.. Mirrored volumes are available only on systems that support RAID, see `sys raid`.

device-agent

Displays the name of the service which is responsible for managing the type of block device on which a give image is available. `vcmp-virtual-cdrom` is the device-agent which allows VCMP guests to install images that reside in the host system, via a virtual cdrom device.

file-size

Displays the size of the block-device-image file in megabytes.

glob Displays the items that match the glob expression. See help `glob` for a description of glob expression syntax.

last-modified

Displays the date the file was last modified.

name Specifies the name of the block-device-image that you want to install or delete.

product

Displays the F5 Networks product the block-device-image contains.

reboot

Specifies that the system reboots immediately after a successful installation.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

resource-id

Displays the resource-id string corresponding to software image in question. This string is used to identify the image for use in interacting with the device agent that is responsible for management of the resource. In the case of the `vcmp-virtual-cdrom` device-agent the resource-id is used by the live installation daemon when requesting that the hypervisor make images available, before the installation can proceed.

verified

When set to `yes`, indicates that the block-device-image is authentic.

version

Displays the version number of the product this block-device-image contains.

volume

Specifies the name of the volume on which you want to install the block-device-image, or from which you want to delete the block-device-image.

Note: You cannot install software on the active volume.

SEE ALSO

`delete`, `glob`, `install`, `list`, `reboot`, `regex`, `sys software block-device-hotfix`, `sys software hotfix`, `sys software image`, `tmsh`, `show`, `sys software status`, `sys version`, `sys license`, `sys provision`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2013-2016. All rights reserved.

BIG-IP 2016-03-14 `sys software block-device-image(1)`

sys software hotfix

NAME

`hotfix` - Manages F5 Networks software hotfixes.

MODULE

`sys software`

SYNTAX

Install, display information about, or delete a hotfix using the syntax in the following sections.

INSTALL

install hotfix [name] volume [name]

options:

create-volume
reboot

DISPLAY

list hotfix

list hotfix [[[name [/slot_id]] | [glob] | [regex]] ...]

options:

build
checksum
id
one-line
product
title
verified
version

DELETE

delete hotfix [[name] ...]

options:

all

DESCRIPTION

You can use the hotfix component to install a hotfix onto a volume, view information about available hotfixes, or delete unwanted hotfixes.

Use the create-volume option with the hotfix component to create new volumes.

Note: You use the slot_id option only for chassis systems and only when displaying the values for the options of a specific hotfix. You do not use the slot_id option when installing or deleting a hotfix, because these commands operate on all blades or the entire system.

EXAMPLES

list hotfix Hotfix-BIGIP-9.6.1-824.0-HF3.im

Displays information about the specified hotfix, BIGIP-9.6.1-824.0-HF3.im.

list hotfix */1

Displays information about the all the hotfixes on the first slot.

install hotfix Hotfix-BIGIP-9.6.1-824.0-HF3.im volume HD1.1 reboot

Attempts to install the specified hotfix, BIGIP-9.6.1-824.0-HF3.im, onto HD1.1.

Note: If the installation is successful, and you used the reboot option, as in this example, the machine reboots into the newly installed hotfix.

OPTIONS

build

Displays the build number of the hotfix.

checksum

Displays the checksum of the hotfix. You can use this option to verify the integrity of the hotfix.

create-volume

Create a new volume using the name specified with the volume option. Mirrored volume names must begin with the prefix MD1.. Mirrored volumes are available only on systems that support RAID, see sys raid.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies the name and sequential ID of the hotfix that you want to install or delete.

product

Displays the F5 Networks product this hotfix contains.

reboot

Specifies that the system reboots immediately after a successful installation.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

slot_id

Specifies the number of the slot on a chassis system that contains the hotfix about which you want to display information.

title

Displays a textual description of the hotfix.

verified
When set to yes, indicates that the hotfix is authentic.

version
Displays the version number of the product the hotfix contains.

volume
Specifies the name of the volume on which you want to install the hotfix, or from which you want to delete the hotfix.

SEE ALSO
delete, glob, install, list, regex, sys software image, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-04-16 sys software hotfix(1)

sys software image

NAME
image - Manages F5 Networks software images.

MODULE
sys software

SYNTAX
Install, display information about, or delete a software image using the syntax in the following sections.

INSTALL
install
options:
create-volume
image [name]
reboot
volume [name]

DISPLAY
list image
list image [[[name [/slot_id]] | [glob] | [regex]] ...]
options:
build
build-date
checksum
file-size
last-modified
one-line
product
verified
version

DELETE
delete image [[[name] ...] | [all]]

DESCRIPTION
You can use the image component to install images onto a volume, view information about available images, or delete unwanted images.

INSTALLING A SOFTWARE IMAGE

Before you begin installing an image, you must download the image file into the /shared/images directory. You can find new software images at <http://downloads.f5.com>. We recommend downloading both the .iso file and the .md5 file. Download the file (or files) to your local machine, then transfer it to the /shared/images directory on the BIG-IP(r). Use the Manager (GUI) interface to make this transfer, or quit tmsh to the Unix command line and use scp or a similar Unix command.

If you downloaded the .md5 file, you can use the Unix md5sum command to check the MD5 hash of the .iso file, and you can compare it to the contents of the .md5 file. They should match. If they do not, retry the download and/or transfer of the .iso file.

From tmsh, you can use show sys software status to see all of the available disk volumes where you can install the .iso file. You can install the .iso file in any volume that is not active.

Then use the install command with this component to install the .iso file to an unused volume. You can use the create-volume option if you want to create a new volume. The installation takes some time; you can use show sys software status repetitively to watch the progress of the installation. To put the .iso file into active service, use the reboot option in the install command, or use the reboot volume vol-name command after the install command completes.

Note: You use the slot_id option only for chassis systems and only when displaying the values for the options of a specific image. You do not use the slot_id option when installing or deleting an image, because these commands operate on all blades or the entire system.

CONFIRMING AN IMAGE INSTALLATION

You can use show sys version to confirm that the system is running the new software version. If this is a new module for the current system, you may need to use show sys license and/or install sys license to update your license. For a new module, you may also need to provision CPU, memory, and disk space for the module with the sys provision component.

EXAMPLES

```
install image BIGIP-10.0.0.5376.0.iso volume HD1.1 reboot
```

Attempts to install the specified image, BIGIP-10.0.0.5376.0.iso, onto HD1.1. **Note:** If the installation is successful and the version is permitted, the machine reboots into the newly installed image. If the version is not permitted, add forced to the command to force the reboot.

```
list image BIGIP-10.0.0.5376.0.iso
```

Displays information about the specified image, build 5376.0 of BIG-IP version 10.0.0.

```
list image */1
```

Displays information about all of the images located on the first slot.

OPTIONS

build

Displays the build number of the image.

build-date

Displays the date on which the image was built.

checksum

Displays the checksum of the image. You can use this option to verify the integrity of the image.

create-volume

Creates a new volume using the name specified with the volume option. Mirrored volume names must begin with the prefix MD1.. Mirrored volumes are available only on systems that support RAID, see sys raid.

file-size

Displays the size of the image file in megabytes.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

last-modified

Displays the date the file was last modified.

name Specifies the name of the image that you want to install or delete.

product

Displays the F5 Networks product the image contains.

reboot

Specifies that the system reboots immediately after a successful installation.

forced

Forces a reboot if the version is not permitted by the license.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

verified

When set to yes, indicates that the image is authentic.

version

Displays the version number of the product this image contains.

volume

Specifies the name of the volume on which you want to install the image, or from which you want to delete the image.

Note: You cannot install software on the active volume.

SEE ALSO

delete, glob, install, list, reboot, regex, sys software hotfix, tmsh, show, sys software status, sys version, sys license, sys provision

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2019-05-02 sys software image(1)

sys software signature

NAME

signature - Manages F5 Networks software signatures.

MODULE

sys software

SYNTAX

Display information about, or delete a signature using the syntax in the following sections.

DISPLAY

list signature

list signature [[[name [/slot_id]] | [glob] | [regex]] ...]

options:

one-line

DELETE

delete signature [[name] ...]

options:

all

DESCRIPTION

You can use the signature component to view information about available signatures, or delete unwanted signatures.

Note: You use the `slot_id` option only for chassis systems and only when displaying the values for the options of a specific signature. You do not use the `slot_id` option when deleting a signature, because these commands operate on all blades or the entire system.

EXAMPLES

```
list signature BIGIP-11.5.0.0.135.iso.sig
```

Displays information about the specified signature, `BIGIP-11.5.0.0.135.iso.sig`.

```
list signature */1
```

Displays information about the all the signatures on the first slot.

OPTIONS

regex

Displays the items that match the regular expression. The regular expression must be preceded by an `@` (`@[regular expression]`) to indicate that the identifier is a regular expression. See help `regex` for a description of regular expression syntax.

slot_id

Specifies the number of the slot on a chassis system that contains the hotfix about which you want to display information.

SEE ALSO

delete, glob, list, regex, sys software image, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-11-18 sys software signature(1)

sys software status

NAME

status - Displays the status of a BIG-IP(r) system software installation.

MODULE

sys software

SYNTAX

Display information about the status component within the sys software module using the following syntax.

DISPLAY

show status

options:

field-fmt

DESCRIPTION

You can use the status component to display the status of the software installation, including whether the system is active, the name of the product being installed, the software version and build number of the software, and the slot and volume on which the software is installed.

After you use the install sys software image command (see install and "sys software image") to install a new software image, you can use this command to monitor the progress of the installation. A percentage meter appears in the Status column.

EXAMPLES

show status

Displays the status of the software installation in a table.

show status field-fmt

Displays the status of the software installation separately for each volume on the system.

```
root@(big-ip1)(cfg-sync Standalone)(Active)(/Common)(tmsh)# quit
[root@big-ip1:Active:Standalone] images # watch tmsh show sys software status
```

Launches the Unix watch command from the Unix command line. The command produces auto-updating output similar to this:

```
Every 2.0s: tmsh show sys software status Thu Oct 18 14:04:04 2012
```

```
-----
Sys::Software Status
Volume Product Version Build Active Status
-----
HD1.1 EM 3.2.0 222.0 no installing 6.000 pct
HD1.2 EM 3.2.0 150.0.465 yes complete
HD1.3 EM 3.2.0 67.0 no complete
```

Where the "installing 6.000 pct" status increases until it eventually changes to "complete." It changes to a specific failure message if there is an issue.

OPTIONS

field-fmt

Specifies to display the software status for each volume in a field format, rather than in a table.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys software status(1)

sys software update-status

NAME

update-status - Displays the BIG-IP(r) update check results.

MODULE

sys software

SYNTAX

Display the results of an update check contained in the update-status component within the sys software module using the syntax in the following section.

DISPLAY

list update-status

options:

all-properties

one-line

[update type] (e.g. RELEASE)

show update-status

options:

field-fmt

[update type] (e.g. RELEASE)

DESCRIPTION

You can use the update-status component to display the results of the update check feature.

EXAMPLES

list update-status

Displays all update check information for the system.

show update-status RELEASE

Displays update check information for the RELEASE update type for the system in a formatted output.

list update-status GEOLOC all-properties one-line

Displays all update check information for the GEOLOC update type on one line.

list update-status last-checked-version

Displays the last checked version for all update types.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

available

This is the file name of the available update.

check-user

This is the system user that last executed the update check.

label

This is the label used when displaying the status on the GUI.

last-checked

This is the last time this update type was checked.

last-checked-auto-mode

This is false if the last time this update type was checked was performed manually.

last-checked-version

This is the version found at the last time this update type was checked.

progress-status

This is the state of the update check.

supplement

This is the file name of the supplemental file.

url This is the URL linking to the available update.

url-supplement

This is a URL linking to a file supplemental to the available update.

For information about the options that you can use with the command list, see help list.

SEE ALSO

list, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2015. All rights reserved.

sys software update

NAME

update - Displays the BIG-IP(r) update check and phone home schedule settings.

MODULE

sys software

SYNTAX

Display and modify the update component within the sys software module using the syntax in the following section.

MODIFY

modify update

options:

auto-check

auto-phonehome

frequency

DISPLAY

list update

options:

all-properties

one-line

show update

options:

field-fmt

DESCRIPTION

You can use the update component to display or modify the configuration of the update check and phone home feature.

EXAMPLES

list update

Displays update check and phone home configuration information for the system.

show update

Displays update check and phone home configuration information for the system formatted for easy viewing.

modify update frequency monthly

Modify the frequency of update checks to monthly.

modify update auto-check disabled

Disable the auto update check feature.

modify update auto-phonehome disabled

Disable the auto phone home feature.

OPTIONS

auto-check

Set this to enabled in order to turn on the auto update check feature. disabled turns the feature off.

auto-phonehome

Set this to enabled in order to turn on the auto phone home feature. disabled turns the feature off.

check-status

This read-only field displays the result of the last update check.

errors

This read-only field displays the number of consecutive errors detected by update checking.

frequency

The frequency of update checks can be one of daily, weekly, or monthly.

For information about the options that you can use with the command list, see help list.

SEE ALSO

list, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010, 2013, 2015-2016. All rights reserved.

BIG-IP 2016-03-14 sys software update(1)

sys software volume

NAME

volume - Manages software volumes on the BIG-IP(r) system.

MODULE

sys software

SYNTAX

Delete, reboot into, or display information about a hard drive volume using the syntax in the following sections.

REBOOT

reboot volume [name]

DISPLAY

list volume

list volume [[[name].[slot_id]] | [glob] | [regex]] ...]

show running-config

show running-config [[[name].[slot_id]] | [glob] | [regex]] ...]

options:

active

active-requested

all-properties

basebuild

build

edition

media [media] [size] [default-boot-location]

one-line

product

status

version

DELETE

delete volume [name]

DESCRIPTION

You can use the volume component to view information about configured volumes, delete unwanted volumes, and reboot the device to a specific volume.

Volumes are created using the install command. See help sys software image and the option create-volume.

Deleting or rebooting into a volume on a VIPRION system affects the entire chassis; therefore, you do not need to specify the slot number.

EXAMPLES

list volume */1

Displays the details of all the volumes located on the first slot in a chassis.

delete volume HD1.5

Deletes the volume named HD1.5.

reboot volume HD1.1

Boots into volume HD1.1 if that volume is not already active. If the volume has an image actively being installed on it, the reboot occurs when the installation is complete.

OPTIONS

active

Specifies if this volume is being run.

active-requested

Specifies if this volume should be active once its status is complete. The system associates this setting

with either the active volume or the volume that is going to become active when its status is complete. If active-requested is set on a volume that is not presently active, the system reboots into the volume when the volume status is complete. As an example, install sys software image BIGIP-10.1.0.3341.0.iso volume HD1.2 reboot will cause active-requested to be set on volume HD1.2, and the system will reboot into volume HD1.2 when the installation is complete. This value is read-only.

basebuild

Displays the build number of either the hotfix presently applied to the system or the original build.

build

Displays the original build number (before any hotfixes).

edition

Displays a textual description of the image. You can use this option to specify the hotfix you want to install.

media

Displays a description of the physical media on which the volume exists. The options are:

media

The type of physical device on which the volume exists, for example, hard drive (hd) or compact flash (cf).

size The space on the slot reserved for the volume.

default-boot-location

Specifies the volume into which the system boots if the slot resets.

name Specifies the name of the volume you are configuring. Volume names are in the format HDX.Y, CFX.Y, or MDX.Y, where X is the hard drive index (HDX), compact flash index (CFX), or RAID index (MDX) (on systems that support RAID), and Y is the volume number on that drive.

product

Displays the F5 Networks product that is installed on the volume.

reboot

Reboots the system into the specified volume.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

status

Displays the installation status of the volume. The options are complete or installing.

version

Displays the version number of the software installed on the volume.

SEE ALSO

delete, glob, install, list, reboot, regex, sys software hotfix, sys software image, sys raid, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2013-03-21 sys software volume(1)

sys sshd

NAME

sshd - Configures the Secure Shell (SSH) daemon for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Configure the sshd component within the sys module using the syntax in the following sections.

MODIFY

modify sshd

options:

allow [add | delete | replace-all-with] {

```
[ [hostname] | [IP address] ] ...
}
allow none
banner [disabled | enabled]
banner-text [string]
inactivity-timeout [integer]
include [string]
login [disabled | enabled]
log-level [debug | debug1 | debug2 | debug3 | error | fatal |
info | quiet | verbose]
port [integer]
```

edit sshd
options:
all-properties
non-default-properties

DISPLAY
list sshd
list sshd [option]
show running-config sshd
show running-config sshd [option]
options:
all-properties
non-default-properties
one-line

DESCRIPTION

You can use the sshd component to configure a secure channel between the BIG-IP system and other devices.

F5 Networks recommends that users of the Configuration utility exit the utility before changes are made to the system using the sshd component. This is because making changes to the system using this component causes a restart of the sshd daemon. Likewise, restarting the sshd daemon creates the necessity for a restart of the Configuration utility.

EXAMPLES

```
modify sshd allow add {192.168.0.0/255.255.0.0}
```

Creates an initial range of IP addresses (192.168.0.0 with a netmask of 255.255.0.0) that are allowed to log in to the system.

```
modify sshd allow add {192.168.1.245}
```

Adds the IP address, 192.168.1.245, to the existing list of IP addresses that are allowed to log in to the system.

```
modify sshd login enabled
```

Enables SSH login to the system.

```
modify sshd inactivity-timeout 3600
```

Sets an inactivity timeout of 60 minutes for SSH logins to the system.

```
modify sshd log-level error
```

Sets the sshd message log level to ERROR.

```
modify sshd banner enabled banner-text "NOTICE: Improper use of this computer may result in prosecution!"
```

Creates a banner that displays when a user attempts to log in to a system using SSH.

Note that you must enclose the banner text in double quotation marks, and then type single quotation marks outside the double quotation marks. You can also use the backslash character to escape each quotation mark as well as any other special characters that the system might process (for example, exclamation point !).

OPTIONS

allow
Configures servers in the /etc/hosts.allow file. The default value is all.

Warning: Using the value none resets the sshd daemon to allow all servers access to the system. F5 Networks recommends that you do not use the value none with the sshd component.

banner
Enables or disables the display of the banner text field when a user logs in to the system using SSH. The default value is disabled.

banner-text
When the banner option is enabled, specifies the text to include in the banner that displays when a user attempts to log on to the system.

fips-cipher-version
Read-only field for internal use. Non-zero value indicates that the list of ciphers has been set to FIPS 140-2 compliant defaults. The value 1 indicates that the list of ciphers is "aes128-cbc,aes256-cbc". User changes to the list of ciphers will not affect the value of this field. This field is relevant only when FIPS 140-2 compliance is enabled in the license.

inactivity-timeout

Specifies the number of seconds before inactivity causes an SSH session to log out. The default value is 0 (zero) seconds, which indicates that inactivity timeout is disabled.

include

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using the include option. If you use this option incorrectly, you put the functionality of the system at risk.

login

Enables or disables SSH logins to the system. The default value is enabled.

log-level

Specifies the minimum sshd message level to include in the system log. The possible values are:

debug - debug3

Indicates that the minimum sshd message level that the system logs is the specified debugging level of messages.

error

Indicates that the minimum sshd message level that the system logs is error.

fatal

Indicates that the minimum sshd message level that the system logs is fatal.

info Indicates that the minimum sshd message level that the system logs is informational.

quiet

Indicates that the system does not log sshd messages.

verbose

Indicates that the system logs all sshd messages.

port Specifies the TCP port to run SSHD. It is a number in the range of 1 and 65535.

The default value is 22.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010, 2012-2013, 2016. All rights reserved.

BIG-IP 2017-09-07 **sys sshd(1)**

sys state-mirroring

NAME

state-mirroring - Configures connection mirroring for a BIG-IP(r) system that is part of a redundant pair in a high availability system.

MODULE

sys

SYNTAX

Configure the state-mirroring component within the sys module using the syntax in the following sections.

MODIFY

modify state-mirroring

options:

addr [IP address]
peer-addr [IP address]
secondary-addr [IP address]
secondary-peer-addr [IP address]
state [enabled | disabled]

edit state-mirroring

options:

all-properties
non-default-properties

DISPLAY

list state-mirroring
list state-mirroring [option]
show running-config state-mirroring
show running-config state-mirroring [option]
options:
all-properties
non-default-properties
one-line

DESCRIPTION

You can use this component to configure connection mirroring on a system that is part of a redundant pair in a high availability system.

Connection mirroring is the process of duplicating connections from the active system to the standby system. Enabling this setting ensures a higher level of connection reliability, but it may also have an impact on system performance.

EXAMPLES

```
modify state-mirroring state enabled addr 192.168.10.10 peer-addr 192.168.10.20
```

Enables and configures connection mirroring for a high availability system in which one BIG-IP system has an IP address of 192.168.10.10, and its peer has an IP address of 192.168.10.20.

```
modify state-mirroring state enabled
```

Re-enables connection mirroring for a system for which connection mirroring was disabled.

OPTIONS

addr Specifies the primary self-IP address on this unit to which the peer unit in this redundant pair mirrors its connections. The default value is ::.

peer-addr
Specifies the primary self-IP address on the peer unit to which this unit mirrors its connections. The default value is ::.

secondary-addr
Specifies another self-IP address on this unit to which the peer unit mirrors its connections when the primary address is unavailable. The default value is ::.

secondary-peer-addr
Specifies another self-IP address on the peer unit to which this unit mirrors its connections when the primary peer address is unavailable. The default value is ::.

state
Enables or disables connection mirroring. The default value is enabled.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2010. All rights reserved.

BIG-IP 2010-03-31 sys state-mirroring(1)

sys sync-sys-files

NAME

sync-sys-files - This command has been removed.

MODULE

sys

SYNTAX

This command has been removed.

RUN

```
run sync-sys-files
```

options:

```
from [IP address]
```

DISPLAY

```
show sync-sys-files
```

DESCRIPTION

This command has been removed.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013. All rights reserved.

BIG-IP 2017-02-23 sys sync-sys-files(1)

sys syslog

NAME

syslog - Configures the BIG-IP(r) system log.

MODULE

sys

SYNTAX

Configure the syslog component within the sys module using the syntax in the following sections.

MODIFY

modify syslog

options:

auth-priv-from [alert | crit | debug | emerg | err | info | notice | warning]
auth-priv-to [alert | crit | debug | emerg | err | info | notice | warning]
clustered-host-name [enabled | disabled]
clustered-message-slot [enabled | disabled]
cron-from [alert | crit | debug | emerg | err | info | notice | warning]
cron-to [alert | crit | debug | emerg | err | info | notice | warning]
daemon-from [alert | crit | debug | emerg | err | info | notice | warning]
daemon-to [alert | crit | debug | emerg | err | info | notice | warning]
description [string]
include [string]
iso-date [enabled | disabled]
console-log [enabled | disabled]
kern-from [alert | crit | debug | emerg | err | info | notice | warning]
kern-to [alert | crit | debug | emerg | err | info | notice | warning]
local6-from [alert | crit | debug | emerg | err | info | notice | warning]
local6-to [alert | crit | debug | emerg | err | info | notice | warning]
mail-from [alert | crit | debug | emerg | err | info | notice | warning]
mail-to [alert | crit | debug | emerg | err | info | notice | warning]
messages-from [alert | crit | debug | emerg | err | info | notice | warning]
messages-to [alert | crit | debug | emerg | err | info | notice | warning]
remote-servers [add | delete | modify | replace-all-with] {
 [name] {
options:
 host [hostname]
 local-ip [IP address]
 remote-port [port number]
 }
 }
remote-servers none
user-log-from [alert | crit | debug | emerg | err | info | notice | warning]
user-log-to [alert | crit | debug | emerg | err | info | notice | warning]

edit syslog

options:
all-properties
non-default-properties

DISPLAY

list syslog
list syslog [option]
show running-config syslog
show running-config syslog [option]
options:
all-properties
non-default-properties
one-line

DESCRIPTION

You can use the syslog component to configure the system log.

EXAMPLES

modify syslog auth-priv-from warning

Resets the lowest level of messages about user authentication that are included in the system log to messages with a level of warning, error, critical, alert, and emergency.

modify syslog auth-priv-to warning

Resets the highest level of messages about user authentication that are included in the system log to messages with a level of warning, error, critical, alert, and emergency.

OPTIONS

auth-priv-from

Specifies the lowest level of messages about user authentication to include in the system log. The default value is notice.

auth-priv-to

Specifies the highest level of messages about user authentication to include in the system log. The default value is emerg.

clustered-host-name

When enabled, the slash-separated slot ID of the blade that originated the log is prepended to the hostname field. The default value is enabled.

clustered-message-slot

When enabled, the space-separated slot ID of the blade that originated the log is prepended to the message text. The default value is disabled.

cron-from

Specifies the lowest level of messages about time-based scheduling to include in the system log. The default value is warning.

cron-to

Specifies the highest level of messages about time-based scheduling to include in the system log. The default value is emerg.

daemon-from

Specifies the lowest level of messages about daemon performance to include in the system log. The default value is notice.

daemon-to

Specifies the highest level of messages about daemon performance to include in the system log. The default value is emerg.

description

User defined description.

host Specifies the IP address of a remote server to which syslog sends messages. The default value is none.

include

Warning: Do not use this option without assistance from the F5 Technical Support team. The system does not validate the commands issued using the include options. If you use this option incorrectly, you put the functionality of the system at risk.

iso-date

Enables or disables the ISO date format for messages in the log files. The default value is disabled.

console-log

Enables or disables logging emergency syslog messages to the console. The default value is enabled.

kern-from

Specifies the lowest level of kernel messages to include in the system log. The default value is debug.

kern-to

Specifies the highest level of kernel messages to include in the system log. The default value is emerg.

local-ip

Specifies the IP address of the interface syslog binds with in order to log messages to a remote host. For example, if you want syslog to log messages to a remote host that is connected to a VLAN, you set this parameter to the self IP address of the VLAN.

local6-from

Specifies the lowest error level for messages from the local6 facility to include in the log. The default value is notice.

local6-to

Specifies the highest error level for messages from the local6 facility to include in the log. The default value is emerg.

mail-from

Specifies the lowest level of mail log messages to include in the system log. The default value is notice.

mail-to

Specifies the highest level of mail log messages to include in the system log. The default value is emerg.

messages-from

Specifies the lowest level of messages about user authentication to include in the system log. The default value is notice.

messages-to

Specifies the highest level of system messages to include in the system log. The default value is warning.

remote-port

Specifies the port number of a remote server to which syslog sends messages. The default value is 514.

remote-servers

Configures the remote servers, identified by IP address, to which syslog sends messages. The default value is none.

user-log-from

Specifies the lowest level of user account messages to include in the system log. The default value is notice.

user-log-to

Specifies the highest level of user account messages to include in the system log. The default value is emerg.

SEE ALSO

edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2018-06-29 sys syslog(1)

sys tmm-info

NAME

tmm-info - Displays information about the Traffic Management Microkernel (tmm) daemon.

MODULE

sys

SYNTAX

Display statistics for the tmm-info component within the sys module using the syntax in the following section.

DISPLAY

show tmm-info

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

global

DESCRIPTION

You can use the tmm-info component to display information about the tmm daemon. The purpose of this daemon is to direct all application traffic passing through the BIG-IP(r) system.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, sys tmm-traffic, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 sys tmm-info(1)

sys tmm-traffic

NAME

tmm-traffic - Displays Traffic Management Microkernel (tmm) statistics.

MODULE

sys

SYNTAX

Configure the tmm-traffic component within the sys module using the syntax in the following section.

MODIFY

reset-stats tmm-traffic

DISPLAY

show tmm-traffic

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
global

DESCRIPTION

You can use the tmm-traffic component to display tmm traffic statistics, including errors and redirected connections. The purpose of this daemon is to direct all application traffic passing through the BIG-IP(r) system.

OPTIONS

For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO

reset-stats, show, sys tmm-info, sys traffic, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 sys tmm-traffic(1)

sys traffic

NAME

traffic - Displays or resets traffic statistics for the system.

MODULE

sys

SYNTAX

Configure the traffic component within the sys module using the syntax in the following section.

MODIFY
reset-stats traffic

DISPLAY
show traffic
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION
You can use the traffic component to display traffic statistics, including for client, server, Packet Velocity(r) ASIC (PVA), miscellaneous, and authorization traffic. You can also reset the traffic statistics to zero at any time.

OPTIONS
For information about the options that you can use with the command show, see help show.

For information about the options that you can use with the command reset-stats, see help reset-stats.

SEE ALSO
reset-stats, show, sys pva-traffic, sys tmm-info, sys tmm-traffic, tmsh

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2011-04-25 sys traffic(1)

sys turboflex features

NAME
features - Displays all TurboFlex features.

MODULE
sys turboflex

SYNTAX
Displays current features component within the sys turboflex module using the syntax in the following section.

DISPLAY
show features
options:
all-properties
field-fmt
current-module

DESCRIPTION
You can use the features component to display the active TurboFlex profiles and its feature set.

EXAMPLES
show features

Displays the active TurboFlex profiles and its feature set.

SEE ALSO
show, tmsh, turboflex

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-03 sys turboflex features(1)

sys turboflex profile-config

NAME
profile-config - Configures the TurboFlex profile chosen by users.

MODULE
sys turboflex

SYNTAX
Configure the profile-config component within the sys turboflex module using the syntax shown in the following sections.

MODIFY
modify profile-config
options:
type [turboflex-adc | turboflex-highspeed-layer4 | turboflex-security | turboflex-base | turboflex-low-latency | turboflex-dn:

DISPLAY
list profile-config
options:
all-properties

DESCRIPTION
You can use the profile-config component to configure the TurboFlex profile type to use.

EXAMPLES
list profile-config

Displays properties of the current TurboFlex profile configuration.

modify profile-config type

Modify the type of the active TurboFlex profile configuration. The default is turboflex-adc.

OPTIONS
type The type for TurboFlex profile currently chosen on the system.

turboflex-base: The Standard basic profile.

turboflex-adc: The Application Delivery Controller (ADC) profile provides the HW acceleration features normally associated with the ADC use cases. In order to enable this profile, the LTM module must be provisioned.

turboflex-ultrafast-layer4: The Ultra Fast L4 CPS profile provides HW acceleration features normally associated with fast L4 LBing. In order to enable this profile, the LTM module must be provisioned. For full capability, the TAM module must be provisioned after selecting this profile.

turboflex-security: The Security profile provides the HW acceleration features normally associated with the security use cases. In order to enable this profile, the LTM modules must be provisioned. For full capability, the AFM module must be provisioned after selecting this profile.

turboflex-low-latency: The Low Latency FIX profile provides the HW acceleration features normally associated the FIX use case. In order to enable this profile, the Advanced Protocol or FIX-LL add-ons must be licensed and the LTM module must be provisioned.

turboflex-private-cloud: The Private Cloud (PC) profile provides the HW acceleration features normally associated with the PC use cases. In order to enable this profile, the LTM module must be provisioned.

SEE ALSO
list, modify, tmsh, turboflex

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008, 2012-2014, 2017. All rights reserved.

BIG-IP 2017-09-05 sys turboflex profile-config(1)

sys turboflex profile all

NAME

all - Displays each profile and its feature.

MODULE

sys turboflex profile

SYNTAX

Displays current all component within the sys turboflex profile module using the syntax in the following section.

DISPLAY

show all
options:
all-properties
field-fmt
current-module

DESCRIPTION

You can use the all component to display all TurboFlex profiles and their individual feature set.

EXAMPLES

show all

Displays all TurboFlex profiles and their individual feature set.

SEE ALSO

show, tmsh, turboflex

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-09-05 sys turboflex profile all(1)

sys turboflex profile feature

NAME

feature - Displays information for the Active TurboFlex profile.

MODULE

sys turboflex profile

SYNTAX

Displays current feature component within the sys turboflex profile module using the syntax in the following section.

DISPLAY

show feature
options:
all-properties
field-fmt
current-module
running-config

DESCRIPTION

You can use the feature component to display information about current Active TurboFlex profiles.

EXAMPLES

show feature

Displays information on current Active TurboFlex profiles.

SEE ALSO

show, tmsh, turboflex

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-03 sys turboflex profile feature(1)

sys turboflex warning

NAME

warning - Displays all warnings for a TurboFlex profile configuration.

MODULE

sys turboflex

SYNTAX

Displays current warning component within the sys turboflex module using the syntax in the following section.

DISPLAY

show features
options:
all-properties
field-fmt
current-module

DESCRIPTION

You can use the warning component to display the active TurboFlex profiles and its feature set.

EXAMPLES

show warning

Displays the warnings after a TurboFlex profile is configured.

SEE ALSO

show, tmsch, turboflex

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2017. All rights reserved.

BIG-IP 2017-03-29 sys turboflex warning(1)

sys ucs

NAME

ucs - Loads or saves a UCS (.ucs) file.

MODULE

sys

SYNTAX

Configure the ucs component within the sys module using the syntax in the following sections.

MODIFY

save ucs [file name]
options:
no-private-key
passphrase

load ucs [file name]
options:
include-chassis-level-config
no-license
no-platform-check
passphrase
platform-migrate
reset-trust

delete ucs [file name]

DISPLAY

`list ucs`
`show ucs [file name]`

DESCRIPTION

You can use the `ucs` component to save the running configuration of the system into a UCS file. Additionally, you can modify the running configuration of the system by loading an existing UCS file.

When you save a UCS file, the file is saved to the default directory, `/var/local/ucs`.

When you load a UCS file in shell mode, the system searches for the file using the relative path to the default directory (`/var/local/ucs`). When you load a UCS file in bash mode, the system searches the current directory first. If the file is not found in the current directory, the default directory is then searched.

EXAMPLES

`save ucs myucs`

Saves the running configuration of the system into the file `myucs.ucs`.

`load ucs myucs`

Modifies the running configuration of the system by loading the configuration contained in the `myucs.ucs` file.

`delete ucs myucs`

Delete `myucs.ucs` in the default directory, `/var/local/ucs/`.

`list ucs`

Displays existing UCS files in the default directory, `/var/local/ucs/`.

OPTIONS

`include-chassis-level-config`

During restore of the UCS file, include chassis level configuration that is shared among boot volume sets. For example, cluster default configuration.

`no-license`

Performs a full restore of the UCS file and all the files it contains, with the exception of the license file. The option must be used to restore a UCS on RMA devices (Returned Materials Authorization).

`no-platform-check`

Bypasses the platform check and allows a UCS that was created using a different platform to be installed. By default (without this option), a UCS created from a different platform is not allowed to be installed.

`no-private-key`

Indicates that the UCS file can be saved without private key information.

`passphrase`

Specifies the passphrase that is necessary to load the specified UCS file.

`platform-migrate`

Ignore the configuration items specific to a particular platform. The platform-specific objects already present on the device are not removed or modified. The ones included in the UCS archive are not loaded.

This option is not valid for UCS archives from systems provisioned with modules other than LTM and GTM. If the UCS comes from a system running any release of TMOS version 10, only LTM (not GTM) is supported.

This option implies `no-license` and `no-platform-check`.

`reset-trust`

When specified, the device and trust domain certs and keys are not loaded from the UCS. Instead, a new set is regenerated.

SEE ALSO

`load`, `list`, `save`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2014, 2016. All rights reserved.

BIG-IP 2016-03-31 sys ucs(1)

sys url-db download-result

NAME

download-result - Lists download result for URL Master and RTSU DB.

MODULE

sys url-db

SYNTAX

List download-result component within the module using the syntax shown in the following sections.

DISPLAY

The download-result consists of the object name (/Common/masterdb or /Common/rtsudb), and version. These objects are created by BIGIP and cannot be modified or deleted.

list url-db download-result [masterdb | rtsudb]/slot_number

options:

all-properties
non-default-properties
one-line

list url-db download-result masterdb/slot_number

db-version [integer]
ret-code 0

list url-db download-result rtsudb/slot_number

db-version [integer]
ret-code 0

DESCRIPTION

Lists download result for Master URL database and Real-Time Security Update (RTSU). These objects are created after the first successful download and updated after every download. For cluster environments the slot number will be shown after the name of the database.

OPTIONS

db-version

Specifies database version for URL Master or Real-Time Security Update DB.

ret-code

Specifies the download result status and always zero now.

SEE ALSO

sys url-db download-schedule sys url-db url-category

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013-2014. All rights reserved.

BIG-IP 2017-01-19 sys url-db download-result(1)

sys url-db download-schedule

NAME

download-schedule - Configures download schedule for URL Master DB.

MODULE

sys url-db

SYNTAX

Configure a download-schedule component within the module using the syntax shown in the following sections.

MODIFY

The download-schedule consists of the object name (/Common/urlldb), download start time (start-time), download end time (end-time) and status. You can have only one download schedule and the download occurs daily.

modify url-db download-schedule urlldb

start-time [HH::MM]
end-time [HH::MM]
download-now [true | false]
status [true | false]

use-proxy [true | false]

DISPLAY

list url-db download-schedule urldb
options:
all-properties
non-default-properties
one-line

DESCRIPTION

Configures download schedule for Master URL database.

EXAMPLES

```
modify download-schedule urldb { start-time 2:00 end-time 4:00 }
```

Modify the download schedule for Master DB download schedule between 2:00 AM and 4:00 AM. Other downloads such as RTSU (Real-Time Security Update) and ACE (Advanced Classification Engine) DB download occurs at regular intervals.

```
modify download-schedule urldb { start-time 20:00 end-time 22:00 }
```

Modify the download schedule for Master DB download schedule between 8:00 PM and 10:00 AM.

```
modify download-schedule urldb { download-now true }
```

Master DB Download starts in few minutes after issuing this command. The download-now will be set to false after successful download.

```
modify download-schedule urldb { status false }
```

By setting the status flag to false, download (Master and other DB) will not occur any more.

```
modify download-schedule urldb { use-proxy true }
```

DB Download uses proxy configuration after issuing this command.

OPTIONS

download-now
Specifies to start download in few minutes and no need to wait for the scheduled window.

end-time
Shows download end time. Download will start between scheduled start time and end time.

start-time
Shows download start time. Download will start between scheduled start time and end time.

status
Shows the download status is enabled. By turning to false, download will not occur.

use-proxy
Specifies to use proxy configuration for database download.

SEE ALSO

sys url-db download-result sys url-db url-category

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013-2014. All rights reserved.

BIG-IP 2016-02-25 sys url-db download-schedule(1)

sys url-db url-category

NAME

url-category - Configures URL categories for URL classification and filtering

MODULE

sys url-db

SYNTAX

Configure a url-category component within the module using the syntax shown in the following sections.

CREATE/MODIFY

Each url-category consists of the object name (/Common/Business_and_Economy), a display-name ("Business and Economy") which is a more user-friendly category name, and a category number. The hundreds and thousands of URLs under a url-category are stored in a database. You can create your own url-category (custom category) and you can add more URLs to an existing category (recategorization).

create url-db url-category [name]

options:

display-name [string]
description [string]
initial-disposition [integer]
is-security-category [string]
parent-cat-number [integer]
severity-level [integer]
urls [add | delete | modify | replace-all-with] {
[string]
}

modify url-db url-category [name]

initial-disposition [integer]
is-security-category [string]
parent-cat-number [integer]
severity-level [integer]
urls [add | delete | modify | replace-all-with] {
[string]
}

DISPLAY

list url-category

list url-category [[[name] | [glob] | [regex]] ...]

options:

all-properties
non-default-properties
one-line
partition

DESCRIPTION

Configures a url-category

NOTE: When you create a new url-category, you must provide a display-name. However, after creation it cannot be changed to another value. The system will provide a cat-number for your newly created url-category. The number is an integer greater than 1900. The url-category you create is considered to be a custom URL category, and so the is-custom flag will be set to true.

NOTE: The only change you can make to a system provided url-category is to add one or more URLs to its list of URLs. This is called recategorization, and the is-recategory flag will be set to true. You need to do this if the URL does not already exist in the database.

EXAMPLES

```
create url-category my-own-url-cat display-name "My Own URL Category" urls add { http://a.url.com
http://www.another.url.org }
```

Creates a new url-category. The new url-category you create is known as a custom category, as opposed to a system provided url-category. In this case, you must specify the display-name and at least one URL.

```
modify url-category my-own-url initial-disposition 4 parent-category 0
```

Modify the initial-disposition and parent-category in a customized url-category.

```
modify url-category Business_and_Economy urls add { http://www.theneomaxist.com }
```

Modify a system provided url-category by adding a URL to it. This action is called recategorization. The url-category is recategorized.

OPTIONS

cat-number

Shows a unique category number. Custom URL categories have numbers greater than 1900. This is a read-only attribute.

description

Specifies a unique description for the URL category.

display-name

Specifies a user-friendly name that describes what the URL category represents. This attribute cannot be changed after creation.

initial-disposition

Specifies the action to be taken when a certain URL category is not listed in any url-filter.

is-custom

This flag is set by the system when you create your own URL category. This attribute is read-only.

is-security-category

This flag is not being used. This attribute is read-only.

parent-cat-number

Specifies the category number of a parent url-category. 0 denotes no parent.

severity-level
Specifies the severity level.

SEE ALSO

sys url-db download-result sys url-db download-schedule and apm url-filter

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011, 2013, 2016. All rights reserved.

BIG-IP 2016-03-14 sys url-db url-category(1)

sys version

NAME

version - Displays software version information for the BIG-IP(r) system.

MODULE

sys

SYNTAX

Display statistics for the version component within the sys module using the syntax in the following section.

DISPLAY

show version
options:
detail

DESCRIPTION

You can use the version component to display the software version running on the system, including a list of hotfixes that you have applied to the system.

EXAMPLES

show version

Displays software version information.

show version detail

Displays more extensive software version information about the system, including the operating system kernel information, and details about each hotfix that you have applied to the system.

OPTIONS

For information about the options that you can use with the command show, see help show.

SEE ALSO

show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2010. All rights reserved.

BIG-IP 2013-07-26 sys version(1)

util

util ccmode

NAME

ccmode - Set Common Criteria mode and a subset of configuration defaults required for a Common-Criteria-compliant BIGIP system.

MODULE

util

SYNTAX

ccmode

DESCRIPTION

Use this command to set a subset of defaults that are required for a Common-Criteria-compliant system.

Run this command **ONLY** if you are creating a Common-Criteria-compliant BIG-IP system, or have otherwise determined that all of the changes are desirable for your BIG-IP configuration. Note that there are runtime changes triggered by this command.

Running the ccmode command is an essential part of the configuration changes required to configure a BIG-IP system as Common-Criteria-compliant. This command changes the base configuration in several ways, including:

• Defining the minimum required password policy.

• Defining the allowed cipher sets for SSL/TLS.

• Disabling some features excluded from the evaluation and therefore not permitted to be used in a compliant system.

• Setting several DB variables, including the Security.CommonCriteria variable, an indicator from which other runtime changes are triggered.

While running this script is essential to creating a Common-Criteria-compliant system, it is not sufficient. Customers wishing to configure compliant systems must consult the configuration Guidance documentation provided when the evaluation is complete, and follow its instructions to completely configure the BIG-IP.

This command has no facility for "undoing" the changes it makes. Instead, the administrator must reverse or revise all of the individual commands, reset the DB variables to their defaults, save the new configuration, and restart the BIG-IP.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 1996-2015. All rights reserved.

BIG-IP 2017-07-13 util ccmode(1)

util clientssl-ciphers

NAME

clientssl-ciphers - Display the Client SSL ciphers that match a given cipher string.

MODULE

util

SYNTAX

clientssl-ciphers string

DESCRIPTION

Use this command to display all Client SSL ciphers that match the given string.

EXAMPLES

run util clientssl-ciphers default

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015. All rights reserved.

BIG-IP 2015-06-26 util clientssl-ciphers(1)

util diadb

NAME

diadb - Run the diadb command to display, filter or delete dia persistence entries.

MODULE

util

SYNTAX

run util diadb

util dnatutil

NAME

dnat - Command providing reverse and forward mapping for deterministic NAT (DNAT).

MODULE

util

SYNTAX

Run the dnat utility from within the util module using the following syntax:

```
run util dnat [] [
[:]]
```

DESCRIPTION

The dnat utility allows the calculation of forward and reverse source address and port mapping for the deterministic mode of Large Scale NAT.

EXAMPLES

dnat 10.0.0.1 --action forward Shows a list of translation address/port pairs that might be used for a subscriber at 10.0.0.1, using the DNAT states contained in /var/log/itm.

dnat 173.240.102.139:5678 Performs a reverse mapping back to the subscriber address for the connection from 173.240.102.139:5678, using the DNAT states contained in /var/log/itm.

dnat --start_time '2012-09-27 06:30:00' --end_time '2012-09-27 12:10:00' 173.240.102.139:5678 Same as the previous example, but only shows the subscriber addresses that used the translation within the specified time range.

dnat --start_time '2012-09-27 06:30:00' --end_time '2012-09-27 12:10:00' --file /var/log/test 173.240.102.139:5678 Same as the previous example, but use the DNAT states contained in /var/log/test

dnat --file /var/log/test Shows summary information, using the DNAT states contained in /var/log/test

dnat --action summary --start_time '2012-09-27 06:30:00' --end_time '2012-09-27 12:10:00' Shows summary information, using the DNAT states within the specified time range.

OPTIONS

client_addr or --client_addr

Used to provide the subscriber address for forward mappings (--action forward), and the translation address for reverse mappings (--action reverse).

client_port or --client_port

Used to provide the subscriber port for forward mappings (--action forward), and the translation port for reverse mappings (--action reverse).

end_time or --end_time

End time of search range. User can specify the time format via the --time-format switch. The time format defaults to 'YYYY-MM-DD HH:MM:SS.'

start_time or --start_time

Start time of search range. User can specify the time format via the --time-format switch. The time format defaults to 'YYYY-MM-DD HH:MM:SS.'

--action

Specify the action to be taken: summary, forward, reverse, reverse_addr. Default: reverse when supplied

with `client_addr`, summary otherwise.

summary

Provides summary information on the parsed deterministic NAT configuration snippet.

reverse

Returns possible subscriber address for the provided client address and client port (as translation end-point).

forward

Returns possible translation end-points for the provided client address (as subscriber address).

forward_compact

Returns possible subscriber addresses for the provided client address (as translation address), in a compact format with address and port-range.

reverse_addr

Returns possible subscriber addresses for the provided client address (as translation address).

--file /var/log/ltn

Read DNAT state from file (default: /var/log/ltn)

--time_format

Timestamp parse format for `--start_time` or `--end_time` (default: '%F %T %Z', this yields 'YYYY-MM-DD HH:MM:SS', with timezone being optional.)

--flags

DAGLIB flags parameter.

--flags_v2or

DAGLIB flags override parameter, apply when deterministic NAT state is from log version 2 or lower only.

--all

Display all entries, even if there are duplicates (default disabled).

--version

Display version and DAG information.

SEE ALSO

`run`, `tmsh`, `date`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2017-07-13 util dnatutil(1)

util establish adfs trust

NAME

establish-adfs-trust - Establish trust with ADFS server.

MODULE

util

SYNTAX

establish-adfs-trust --vs [--cert] --username [--password]

- Virtual server name. It should have the "adfs_proxy" profile attached.
- Certificate object name for new trust certificate. If the option is not provided, existing trust certificate will be updated.
- / - Credentials of a local administrator account on the ADFS server. The command will prompt for password if it's not provided on the command line.

DESCRIPTION

The establish-adfs-trust utility allows users to establish trust with ADFS server. During trust establishment, APM generates a self signed certificate and registers it with ADFS. On success, APM adds the newly generated certificate and key to the server SSL profile of given virtual server. Any previously attached certificate and key are detached from the server SSL profile, but remain on the system.

EXAMPLES

```
establish-ads-trust --vs ads_vs --cert ads_cert --username Administrator
```

The above command will prompt for password.

```
establish-ads-trust --vs ads_vs --cert ads_cert --username Administrator --password secret
```

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2017. All rights reserved.

BIG-IP 2018-01-11 util establish ads trust(1)

util finalize custom ami

NAME

finalize-custom-ami - Utility to clean-up and prepare the custom AMI for Autoscaling on AWS.

MODULE

util

SYNTAX

Run the finalize-custom-ami utility from within the util module using the following syntax:

```
run util finalize-custom-ami
```

DESCRIPTION

The finalize-custom-ami utility runs the final clean-up steps to prepare the instance for being packaged as a custom AMI that could be autoscaled in AWS. This is a BIG-IP VE on AWS specific utility tool that shouldn't be run on other platforms. This should be used according to supporting documentation provided by F5.

All the output is reported on standard output.

EXAMPLES

```
run util finalize-custom-ami
```

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015. All rights reserved.

BIG-IP 2018-08-22 util finalize custom ami(1)

util geodb

NAME

geodb - Run the geodb command to display Diameter Geo Redundancy status.

MODULE

util

SYNTAX

```
run util diadb
```

util geoutil

NAME

geoutil - Run the geoutil command to change topics, hostnames, and ports the georedundancy config files.

MODULE

util

SYNTAX

run util geoutil

DESCRIPTION

The geoutil utility facilitates changing kafka topics by modifying the configuration files for kafka and mirrormaker. It stops the daemons associated with georedundancy (zookeeper, kafka, and mirrormaker). Then it deletes the old topics. It modifies all the necessary config files. Finally, it restarts the daemons.

Before running geoutil, the remote-host entries in sys global-settings should be modified to reflect the new hostnames.

EXAMPLES

One one machine:

```
modify sys global-settings remote-host modify {
mirrorMakerLocal {
  hostname seattle addr 127.0.0.1 }
mirrorMakerRemote {
  hostname bellevue addr 10.126.2.40 }}
```

```
run util geoutil 10.126.2.1 9092 seattle bellevue
```

On a machine remote from the first:

```
modify sys global-settings remote-host modify {
mirrorMakerLocal {
  hostname bellevue addr 127.0.0.1 }
mirrorMakerRemote {
  hostname seattle addr 10.126.2.30 }}
```

```
run util geoutil 10.126.2.2 9092 bellevue seattle
```

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2019. All rights reserved.

BIG-IP 2019-08-02 util geoutil(1)

util ihealth

NAME

ihealth - Run the ihealth utility for the purpose of uploading a new or existent qkview file to ihealth.

MODULE

util

SYNTAX

```
run util ihealth [ -f ] [ -e ] [ -l ]
[ -a ] [ -u ] [ -p ]
[ -s ] [ -d ] [ -n ] [ -h ] [ -C ]
[ -x ] [ -t ]
[ -z ]
```

or

```
run util ihealth [ --file= ] [ --existing= ]
```

```
[ --auth-url= ] [ --api-url= ]
[ --user= ] [ --password= ] [ --sr= ]
[ --description= ] [ --no-upload ] [ --help ]
[ --exclude-cores ] [ --exclude= ]
[ --upload-timeout ] [ --max-file-size ]
```

DESCRIPTION

The ihealth utility runs a menu that allows users to select qkview files or creating a new qkview file for uploading to ihealth using DevCentral userid.

OPTIONS

--file

This is the name of the local diagnostics file to use or create.

--existing

This flag can be used to specify to use an existing diagnostics file, rather than create one.

--auth-url

This is an over-ride of the DevCentral login server URL used for connecting to the iHealth diagnostic service.

--api-url

This is an over-ride of the iHealth API service used for storing and processing qkviews.

--user

This is an over-ride of the DevCentral user name stored for connecting to the iHealth diagnostic service.

--password

This is an over-ride of the DevCentral password stored for connecting to the iHealth diagnostic service.

--description

This is a description that is attached to the diagnostics file when uploaded to the iHealth diagnostic service.

--sr This is a service request number that is attached to the diagnostics file when uploaded to the iHealth diagnostic service.

--exclude

This is an over-ride of the stored exclude settings for limiting files collected by qkview. See qkview for the possible values.

--exclude-cores

This directs qkview to exclude core files from it's data collection.

--no-upload

This directs the iHealth command to skip uploading to the iHealth service.

--upload-timeout

Enter the maximum number of seconds that the upload of qkview to ihealth can take. The default is zero, meaning that no --max-time will be passed to curl for the upload. Useful for cron jobs.

--max-file-size

Enter the maximum file size for log files that qkview will truncate to (0-100MB). Setting it to zero signifies 100MB. The floor is 2048, passing a value less than this will be treated as 2048.

--help

Displays the usage for this command.

SEE ALSO

run, tmsh, qkview

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2016. All rights reserved.

BIG-IP 2018-05-21 util ihealth(1)

util ipsecalgdb

NAME

ipsecalgdb - Run the ipsecalgdb command to view IPsecALG translation, and pending IKE connection count entries.

MODULE
util

SYNTAX
run util ipsecalgdb
Commands:
delete
list

Objects:
all
tran[slation] entries
pend[ing] IKE connection count entries

DESCRIPTION

The ipsecalgdb utility allows users to view IPsecALG translation, and pending IKE connection count entries, as well as deleting those entries using this utility.

EXAMPLES

run util ipsecalgdb delete all

Delete all IPsecALG translation, and pending IKE connection count entries.

run util ipsecalgdb delete trans

Delete all IPsecALG translation entries.

run util ipsecalgdb delete pend

Delete all IPsecALG pending IKE connection count entries.

run util ipsecalgdb list all

Shows all IPsecALG translation, as well as pending IKE connection count entries.

OPTIONS

delete

Delete all objects of the specified type.

list Display all objects of the specified type.

all = all available object types.

trans = IPsecALG translation entries.

pend = IPsecALG pending IKE connection count entries.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2016-08-02 util ipsecalgdb(1)

util lsndb

NAME

lsndb - Run the lsndb command to view Large Scale NAT persistence entries, inbound mappings, client connection counts, and PCP mappings.

MODULE
util

SYNTAX
run util lsndb
Commands:
del[ete]
list
summary

Objects:
all
client

**inbound[-mapping]
pba
pcp
filters
persist[ence]
all**

DESCRIPTION

The **lsndb** utility allows users to view LSN persistence mappings, inbound mappings, PCP mappings, client connection counts and port block entries. Persistence, inbound and PCP mappings can also be deleted using this utility.

EXAMPLES

run util lsndb delete all

Delete all LSN persistence mappings, inbound mappings and PCP mappings.

run util lsndb del inbound

Delete all LSN inbound mappings.

run util lsndb delete pcp

Delete all PCP mappings.

run util lsndb delete persist

Delete all LSN persistence entries.

run util lsndb list all

Shows all LSN persistence mappings, inbound mappings, PCP mappings, client connection counts and port block entries.

run util lsndb list client

Shows all LSN client connection counts. Each line will display the client IP address and the number of connections used by the client. Connection counts are only available for LSN pools with a non-zero client connection limit.

run util lsndb list inbound

Shows all LSN inbound mappings. Each line will display the translation IP address, the client IP address, the DS-Lite tunnel (if configured) and the age of the mapping.

run util lsndb list pba

Shows all LSN port block entries. Each line will display the client IP address, the port block used and the time that the entry will persist in the database (TTL).

run util lsndb list pcp

Shows all PCP mappings. PCP clients send MAP requests to map their private IP address and port to a public IP address and port. The BIG IP system uses those mappings as NAT entries. Each line will display the client IP address, the external address used and the age of the mapping.

run util lsndb list filters

Shows all LSN filters for inbound mappings. Each line will display the inbound mappings along with filter's remote peer IP address and prefix length.

run util lsndb list persist

Shows all LSN persistence entries. Each line will display the client IP address, the translation address used and the time that the entry will persist in the database (TTL).

run util lsndb summary all

Show summary for all LSN persistence mappings, inbound mappings and port block entries.

run util lsndb summary inbound

Show summary for all LSN inbound mapping entries.

run util lsndb summary pba

Show summary for all LSN port block entries.

run util lsndb summary persist

Show summary for all LSN persistence entries.

OPTIONS

delete

Delete all objects of the specified type. Client connection counts and port block entries cannot be deleted.

list Display all objects of the specified type.

summary

Display summary information of the specified type.

Object types are:

all = all available object types.

client = LSN client counts (list only).

inbound = LSN inbound mapping entries.

pba = LSN port block entries.

pcp = PCP mappings entries.

persist = LSN translation persist entries.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2012-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 util lsndb(1)

util platform check

NAME

platform_check - Runs platform checks available on an active system.

MODULE

util

SYNTAX

Run the platform_check utility from within the util module using the following syntax:

run util platform_check

DESCRIPTION

The platform_check utility runs diagnostics available on an active system to verify correct functionality of platform components. This should be used according to supporting documentation provided by F5.

Output is provided on standard output as well as /var/log/platform_check. Running platform_check with the -h argument will produce available argument listing. Running platform_check with the -l argument will list the available test suites.

EXAMPLES

run util platform_check

Runs all diagnostics available for an active system.

run util platform_check drive

Runs only the drive suite of diagnostics.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2017-07-13 util platform check(1)

util platform diag

NAME

platform_diag - Runs platform diagnostics available on an inactive system

MODULE

util

SYNTAX

Run the platform_diag utility from within the util module using the following syntax:

```
run util platform_diag
```

DESCRIPTION

The platform_diag utility runs the diagnostics on an inactive system to verify correct functionality of platform components. This should not be used on an active system. This should be used according to supporting documentation provided by F5.

Output is provided on standard output as well as /var/log/platform_diag. Running platform_diag with the -h argument will produce available argument listing. Running platform_diag with the -l argument will list the available test suites.

EXAMPLES

```
run util platform_diag
```

Runs all appropriate diagnostics for this platform.

```
run util platform_diag hwaccel
```

Runs only the hwaccel suite of diagnostics.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2016. All rights reserved.

BIG-IP 2016-03-16 util platform diag(1)

util qkcloud

NAME

qkcloud - Run the qkcloud utility for the purpose of displaying cloud meta-data.

MODULE

util

SYNTAX

```
run util qkcloud [ -d ] [ -h ] [ -j ] [ -l ] [ -V ] [ -v ] [ -x ]
```

or

```
run util qkcloud [ --human ] [ --help ] [ --json ] [ --latest ]  
[ --version ] [ --verbose ] [ --xml ]
```

DESCRIPTION

The qkcloud utility displays meta-data obtained from http://169.254.169.254.

OPTIONS

-d or --human

This directs the qkcloud program to output data in human format (similar to yaml).

-h or --help

Displays this help message.

-j or --json

This directs the qkcloud program to output data in json format.

-l or --latest

This tells qkcloud to only show the latest API.

-V or --version

This will show the qkcloud version and the system type.

-v or --verbose

This will tell qkcloud to turn on the CURL verbose mode in order to see the details of communication with http://169.254.169.254.

-x or --xml

This directs the qkcloud program to output data in xml format.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2018. All rights reserved.

BIG-IP 2018-04-20 util qkcloud(1)

util serverssl-ciphers

NAME

serverssl-ciphers - Display the Server SSL ciphers that match a given cipher string.

MODULE

util

SYNTAX

serverssl-ciphers string

DESCRIPTION

Use this command to display all Server SSL ciphers that match the given string.

EXAMPLES

run util serverssl-ciphers default

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2015. All rights reserved.

BIG-IP 2015-06-26 util serverssl-ciphers(1)

util sipdb

NAME

sipdb - Run the sipdb command to view SIP persistence entries.

MODULE

util

SYNTAX

run util sipdb

util ssh keyswap

NAME

ssh-keyswap - Run the ssh-keyswap command to manage SSH keys on the BIG-IP.

MODULE

util

SYNTAX

run util ssh-keyswap

util test-monitor

NAME

test-monitor - Runs an external monitor and displays the inputs to and output from the monitor.

MODULE

util

SYNTAX

Run the test-monitor utility from within the util module using the following syntax:

run util test-monitor address port

DESCRIPTION

The test-monitor utility runs a single instance of a monitor against the specified ip-address:port. The utility output shows the environment, command-line arguments, and resulting messages on stdout and stderr.

Internal monitors are not supported.

EXAMPLES

run util test-monitor monitorA address 10.10.10.4 port 80

Runs a monitor on the IP address 10.10.10.4 and port 80.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2013, 2016. All rights reserved.

BIG-IP 2016-03-14 util test-monitor(1)

util verify encryption

NAME

verify-encryption - Runs the encrypted attributes diagnostics utility

MODULE

util

SYNTAX

Run the verify-encryption utility from within the util module using the following syntax:

```
run util verify-encryption
run util verify-encryption
```

DESCRIPTION

The verify-encryption runs the /usr/sbin/lssa utility that diagnoses issues with encrypted attributes.

When running this utility, it looks through the files /config/bigip.conf, /config/bigip_base.conf, /config/profile_base.conf, /config/bigip_user.conf and returns lines that match the regular expression of keys, but can't actually be decrypted by the master key. Optionally, a file list can be provided as parameters. Only the specified files will be searched in this case.

EXAMPLES

```
run util verify-encryption
```

Runs the lssa utility for the files /config/bigip.conf, /config/bigip_base.conf, /config/profile_base.conf, /config/bigip_user.conf.

```
run util verify-encryption /config/partitions/partA/bigip.conf /config/partitions/partB/bigip.conf
```

Runs the lssa utility for the space separated file-list.

SEE ALSO

run, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2008-2015 All rights reserved.

BIG-IP 2015-10-20 util verify encryption(1)

vcmp

vcmp global

NAME

global - Display global vCMP system statistics.

MODULE

vcmp

SYNTAX

Configure the global component within the vcmp module using the following syntax.

DISPLAY

```
show global
```

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

Use the global component within the vcmp module to display high-level vCMP system statistics on a per-slot basis. These are statistics that are not associated with any particular vCMP guest or virtual-disk.

EXAMPLES

```
show vcmp global
```

Display all global vCMP system statistics.

OPTIONS

For information about the options that you can use with the show command, see help show.

SEE ALSO

tmsh, show, vcmp guest, vcmp virtual-disk

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2012. All rights reserved.

vcmp guest

NAME

guest - Configures a cluster of virtual machines (VMs) that run on one or all slots. This cluster is known as a vCMP guest.

MODULE

vcmp

SYNTAX

Configure the guest component within the vcmp module using the syntax in the following sections.

CREATE

```
create guest [name]
modify guest [name]
options:
  hostname [hostname]
  app-service [[string] | none]
  boot-priority [integer]
  initial-hotfix [hotfix-filename]
  initial-image [image-filename]
  management-gw [ip-address]
  management-ip [ip-address/netmask | ip-address/prefixlen]
  management-network [bridged | isolated]
  slots [integer]
  traffic-profile [vcmp-traffic-profile-name]
  min-slots [integer]
  allowed-slots {
[slot ID] ...
}
  cores-per-slot [integer]
  state [configured | provisioned | deployed]
  virtual-disk [filename]
  vlans [add | delete | replace-all-with] {
[VLAN name] ...
}
  capabilities [add | delete | modify | replace-all-with] {
[capability id] [ { value [integer] } ]
}
}
```

DISPLAY

```
list guest
show guest
```

options:

```
all-properties
status
```

DELETE

```
delete guest [name]
```

DESCRIPTION

Manage vCMP guests running on this host.

EXAMPLES

```
list vcmp guest
```

Lists the current configuration of all guests.

```
show vcmp guest
```

Displays detailed information regarding the state and progress of all guests.

```
show vcmp guest status
```

Displays the running state of all guests, including each guest's prompt status.

```
show vcmp guest all-properties
```

Displays greater detailed statistics and information on all guests.

```
create vcmp guest my_guest slots 4 min-slots 2 management-ip 192.168.45.12/24 management-gw 192.168.45.254
initial-image BIGIP-11.0.0.2400.0.iso
```

Creates a guest that should span four slots, but must span at least two, with the given management IP and gateway, and with the image file BIGIP-11.0.0.2400.0.iso, which is used to install TMOS on the guest's virtual disks. By default, this guest is in the configured state and has a management network in Bridged mode.

`modify vcmp guest my_guest state provisioned`

Moves the guest into the provisioned state, which causes the host to assign the guest to slots, allocate hardware resources to the guest from those slots, and create virtual disks for the guests on those slots.

Moves the guest into the deployed state, which causes the host to start and maintain VMs on each slot that the guest has been assigned to.

`modify vcmp guest my_guest state configured`

Moves the guest back to the configured state, which causes all of its VMs to shut down and the hardware to be deallocated. The guest is unassigned from all slots. The guest's virtual disks will remain on the host.

`modify vcmp guest my_guest traffic-profile fiftyMbpsSLAProfile`

Adds a traffic-profile named fiftyMbpsSLAProfile to the guest in question, which is configured under vcmp traffic-profile.

OPTIONS

`app-service`

Specifies the name of the application service to which the guest belongs. The default value is none. Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the guest. Only the application service can modify or delete the guest.

`boot-priority`

Specifies the boot priority of a guest. Lower values mean higher boot priorities. The default value is 65535. Note: Boot priority is only applied when multiple guests start and hardware resources constrain the number of guests starting.

`hostname`

Assigns the specified host name to the guest. The host name must be a FQDN. If none is given, the default of ".localdomain" is used. If the guest's name contains characters that are not allowed in a FQDN, then "localhost.localdomain" is used.

This is only a suggested value and may be changed on the guest itself. If the guest ever reverts to the default host name, this suggested host name is used instead of the normal system default.

`initial-hotfix`

Specifies which hotfix image to install on newly created virtual disks for this guest. This image is only used when initially creating the virtual disks. After initial creation, the typical live-install process should be used on the guest to manage software upgrades. The image filename must match a verified software image file that exists in the /shared/images directory, otherwise the guest will sit in a wait state on any slot that is missing the hotfix image until that image is added.

This field is required if the guest state is provisioned or deployed, otherwise it can be left blank.

`initial-image`

Specifies which software image to install on newly created virtual disks for this guest. This image is only used when initially creating the virtual disks. After initial creation, the typical live-install process should be used on the guest to manage software upgrades. The image filename must match a verified software image file that exists in the /shared/images directory, otherwise the guest will sit in a wait state on any slot that is missing the software image until that image is added.

This field is required if the guest state is provisioned or deployed, otherwise it can be left blank.

`management-gw`

Specifies the IP address of the default gateway for the management network. This IP address is only a suggested value and can be changed on the guest itself. If the guest ever reverts to the default management gateway, the suggested gateway is used instead of the normal system default.

This field is required if the guest's management-network is bridged, otherwise it can be left blank.

`management-ip`

Specifies the management IP address and netmask to assign to the guest. This address floats to the primary slot of the guest.

This is only a suggested value and can be changed on the guest itself. If the guest ever reverts to the default management IP address, the suggested IP address is used instead of the normal system default.

This field is required if the guest's management-network is bridged, otherwise it can be left blank.

`management-network`

Specifies the management network mode for this guest. When in Bridged mode, the management interfaces on the guest's VMs are bridged to the physical management interfaces on the host blades. This enables the guest to communicate with networks attached to these physical interfaces, the host itself, and other guests in Bridged mode.

In Isolated mode, the management interfaces of the guest's VMs are completely disconnected. The only way to manage such a guest is by connecting to the console on each of the guest's VMs by using the /usr/bin/vconsole utility or by connecting through a configured self IP on a guest's VLAN.

The default value is Bridged.

ssl-mode

Specifies the SSL mode for this guest. When in shared mode the guest shares the available non-dedicated ssl resources with other guests that are in shared mode. when in dedicated mode the guest receives dedicated SSL hardware resources proportional to number of vcpu cores. When in none mode the guest receives no hardware ssl resources. The default value is shared.

slots

Specifies the number of slots to which this guest should be assigned. This number must be greater than zero and no bigger than the cluster size. The host will attempt to assign the guest up to this number of slots.

Note that this property can be changed while the guest is in any state. While in the configured state, modifying the slots property has no effect, since the guest has not yet been assigned to any slots. While in the provisioned state, decreasing this field will cause the guest to be unassigned from enough slots to honor the new value. The host will unassign the guest first from slots that have the most allocated resources. When a guest's slots value is increased, the host attempts to assign the guest to as many slots as possible, up to the new slots value. This same behavior occurs when modifying the property while the guest is in the deployed state, except that running VMs are shut down on any slots that the guest is unassigned from, and new VMs are deployed on any slots to which the guest has been newly assigned.

The default value is 1.

traffic-profile

Specifies a traffic-profile to be used in defining characteristics of traffic which transits the guest's data-plane. For instance a traffic-profile with a color-policer on it that limits the network throughput of the guest may be applied to enforce service agreements between a host admin and a guest user, or to help mitigate network level DOS of other guests in the system.

min-slots

This field dictates the number of slots that the guest must be assigned to. If at the end of any allocation attempt the guest is not assigned to at least this many slots, the attempt fails and the change that initiated it is reverted. A guest's min-slots value cannot be greater than its slots value.

The default value is 1.

allowed-slots

This list contains those slots that the guest is allowed to be assigned to. When the host determines which slots this guest should be assigned to, only slots in this list will be considered. This is a good way to force guests to be assigned only to particular slots, or, by configuring disjoint allowed-slots lists on two guests, that those guests are never assigned to the same slot.

By default this list includes every available slot in the cluster. This means by default the guest is allowed to be assigned to any slot.

cores-per-slot

This value dictates how many cores a guest is allocated from each slot that it is assigned to. Possible values are dependent on the type of blades being used in this cluster. Use tab-completion to see a list of possible values on the current system.

The default cores-per-slot value depends on the type of blades being used in this cluster.

state

Guests are put into the configured state by default. In this state, the configuration for the guest exists on the host, but none of the guest's VMs are running and no hardware resources (for example: CPU cores, memory) are allocated to it. When the guest moves to the provisioned state, hardware resources are allocated to it, and if not already present, virtual disks are created, and the initial-image is installed onto them. In the deployed state, the vcmppd daemon on the host blades use the allocated resources to launch the VMs. Note that moving from the configured state to the deployed state implies the actions that occur in the provisioned state. To shut down a guest's VMs without de-allocating its hardware resources, move the guest from the deployed state to the provisioned state. Moving a guest to the configured state causes its hardware resources to be deallocated. This does not cause the guest's virtual disks to be deleted. They persist on disk and are reused when the vCMP moves back to the provisioned/deployed states.

virtual-disk

Specifies the filename of the virtual disk to use for this guest's VMs. If the filename does not end in .img, it is appended. When the guest moves to a state in which virtual disks need to be provisioned (provisioned or deployed), a new virtual disk image will be created for the guest with this given filename on each slot that the guest is assigned to and does not already have a virtual disk image. The initial-image is used when creating and installing new virtual disk images. If this field is left blank when virtual disk images need to be provisioned for this guest, a default value of ".img" is assigned. If a virtual disk by that name already exists, then an error is thrown. This prevents virtual disks from accidentally being reused by this assigning of default virtual disk filenames.

capabilities

This list contains the various capability flags and an optional value associated with the guest. The possible capability flags are: appliance-mode, stats-isolated-mode, and host-software-only-mode. The value attributes for these capability flags are currently ignored and may be omitted. The capabilities may be added or removed from a vCMP guest in any state.

The appliance-mode capability disables root and bash access to the guest.

The stats-isolated-mode capability prevents some guest statistics from being sent to the hypervisor.

The host-software-only-mode capability prevents the guest from installing images and hotfixes other than

those provided by the hypervisor.

SEE ALSO

create, delete, list, modify, show, tmsh, vcmp global, vcmp virtual-disk

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2013, 2015-2016. All rights reserved.

BIG-IP 2017-04-28 vcmp guest(1)

vcmp health ha-status

NAME

ha-status - Display vCMP guest high availability (HA) status.

MODULE

vcmp health

SYNTAX

Display guest HA status using the following syntax.

DISPLAY

show ha-status

DESCRIPTION

Use the ha-status component within the vcmp health module to display HA status information about the vCMP guests deployed on this system. This is similar to running `tmsh show sys ha-status` inside of a guest.

EXAMPLES

show vcmp health ha-status

Display HA status status for all guests.

show vcmp health ha-status my_guest

Display HA status status for a single guest; "my_guest".

OPTIONS

For information about the options that you can use with the show command, see help show.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2015-11-05 vcmp health ha-status(1)

vcmp health module-provision

NAME

module-provision - Display vCMP guest module provisioning status.

MODULE

vcmp health

SYNTAX

Display guest module provisioning status using the following syntax.

DISPLAY
show module-provision

DESCRIPTION

Use the module-provision component within the vcmp health module to display module provisioning status information about the vCMP guests deployed on this system. This will show you which modules are provisioned in a guest and at what level.

EXAMPLES

show vcmp health module-provision

Display module provisioning status for all guests.

show vcmp health module-provision my_guest

Display module provisioning status for a single guest; "my_guest".

OPTIONS

For information about the options that you can use with the show command, see help show.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2015-11-05 vcmp health module-provision(1)

vcmp health prompt

NAME

prompt - Display vCMP guest prompt status.

MODULE

vcmp health

SYNTAX

Display guest prompt status using the following syntax.

DISPLAY

show prompt

DESCRIPTION

Use the prompt component within the vcmp health module to display the per-slot command-line prompts for the vCMP guests deployed on this system. These are the same prompts that one would see when logging into a guest via SSH. Example: "bigip.mydomain.com:/S1-green-P:Active:Standalone"

EXAMPLES

show vcmp health prompt

Display prompt status for all guests.

show vcmp health prompt my_guest

Display prompt status for a single guest; "my_guest".

OPTIONS

For information about the options that you can use with the show command, see help show.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2015-11-05 vcmp health prompt(1)

vcmp health software

NAME

software - Display vCMP guest software status.

MODULE

vcmp health

SYNTAX

Display guest software status using the following syntax.

DISPLAY

show software

DESCRIPTION

Use the software component within the vcmp health module to display software status information about the vCMP guests deployed on this system. This is similar to running `tmsh show sys software` inside the guest.

EXAMPLES

show vcmp health software

Display software status for all guests.

show vcmp health software my_guest

Display software status for a single guest; "my_guest".

OPTIONS

For information about the options that you can use with the show command, see help show.

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2014. All rights reserved.

BIG-IP 2015-11-05 vcmp health software(1)

vcmp traffic-profile

NAME

traffic-profile - Configures a traffic-profile, which can be applied to a vCMP guest to control characteristics of data-plane network traffic to the guest.

MODULE

vcmp

SYNTAX

Configure the traffic-profile component within the vcmp module using the syntax in the following sections.

CREATE

create traffic-profile [name]

modify traffic-profile [name]

options:

color-policer [color-policer-name]

DISPLAY

list traffic-profile

options:

all-properties

DELETE

delete traffic-profile [name]

DESCRIPTION

Manage vCMP traffic-profiles running on this host.

EXAMPLES

list vcmp traffic-profile

Lists the current configuration of all traffic-profiles.

```
create vcmp traffic-profile fiftyMbpsSLAProfile color-policer fiftyMbpsLimiter
```

Creates a traffic-profile which makes use of the color-policer fiftyMbpsLimiter.

OPTIONS

color-policer

Specifies the color based policer for metering and shaping traffic destined to a guests data-plane.

SEE ALSO

create, delete, edit, glob, list, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2015. All rights reserved.

BIG-IP 2015-05-12 vcmp traffic-profile(1)

vcmp virtual-disk-template

NAME

virtual-disk-template - Manages the vCMP virtual disk templates available on this hypervisor.

MODULE

vcmp

SYNTAX

Configure the virtual-disk-template component within the vcmp module using the syntax in the following sections.

DISPLAY

```
list virtual-disk-template
```

options:

all-properties

DELETE

```
delete virtual-disk-template [name]
```

DESCRIPTION

The virtual-disk-template component is used to list and delete virtual disk templates that are used to create new virtual disk images for vCMP guests. Virtual disk templates are automatically created by vcmpd when guests move to the Provisioned state and a virtual disk template with the version being installed to the guest's virtual disk image doesn't already exist. This is the only way that virtual disk templates are created. Deleting virtual disk templates frees up space but will slow down future virtual disk image installs of that version; the virtual disk template will need to be first re-created.

EXAMPLES

```
list vcmp virtual-disk-template
```

Lists all virtual disk templates currently available.

```
delete vcmp virtual-disk-template my_vdisk
```

Deletes the virtual disk template named my_vdisk. Note that this is only valid if the virtual-disk-template is not currently attached to any vCMP guest.

SEE ALSO

tmsh vcmp virtual-disk

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2014. All rights reserved.

BIG-IP 2015-01-14 vcmp virtual-disk-template(1)

vcmp virtual-disk

NAME

virtual-disk - Manages the vCMP virtual disks available on this hypervisor.

MODULE

vcmp

SYNTAX

Configure the virtual-disk component within the vcmp module using the syntax in the following sections.

DISPLAY

list virtual-disk
show virtual-disk

options:

all-properties

show virtual-disk

DELETE

delete virtual-disk [name]

DESCRIPTION

The virtual-disk component is used to list and delete virtual disks that are used by vCMP guests. Virtual disks are automatically created by vcmpd when guests move to the Provisioned state and do not already have virtual disks attached to them. This is the only way that virtual disks are created. Virtual disks that are not attached to any guest can be deleted. Virtual disks not already in use can be explicitly attached to vCMP guests.

EXAMPLES

list vcmp virtual-disk

Lists all virtual disks currently available.

delete vcmp virtual-disk my_vdisk

Deletes the virtual disk named my_vdisk. Note that this is only valid if the virtual-disk is not currently attached to any vCMP guest.

SEE ALSO

create, delete, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2012. All rights reserved.

BIG-IP 2014-05-08 vcmp virtual-disk(1)

wam

wam ad-policy

NAME

ad-policy - Configures an ad policy for WebAccelerator for use in ad insertion.

MODULE

wam

SYNTAX

Configure the ad-policy within the wam module using the syntax shown in the following sections.

CREATE/MODIFY

```
create ad-policy [name]
modify ad-policy [name]
options:
  ad-insertion-order [random | sequential]
  ads [add | delete | modify] {
    [name] {
options:
  url [url]
  preroll [yes | no]
  }
}
description [string]
```

DISPLAY

```
list ad-policy [name ...]
```

DELETE

```
delete ad-policy [name ...]
```

DESCRIPTION

You can use the ad-policy component to manage the WebAccelerator ad policies. An ad policy defines how the ad insertion is to be performed while processing video resources. Individual ad urls can be configured in the ad-policy along with the insertion order.

EXAMPLES

```
create wam ad-policy my_ad_policy ads add { my_ad1 { preroll yes url http://www.example.com/ad1.m3u8 } }
```

Creates an ad policy named my_ad_policy with an ad named my_ad1 for the url http://www.example.com/ad1.m3u8 and as a preroll candidate.

```
list wam ad-policy my_ad_policy
```

Displays properties of the ad policy named my_ad_policy.

```
delete wam ad-policy my_ad_policy
```

Deletes the ad policy named my_ad_policy.

OPTIONS

ad-insertion-order

Specifies whether the ads are to be inserted randomly or in the order specified in the policy.

ads Specifies the collection of ads.

description

User defined description of an ad policy.

AD OPTIONS

url Specifies the url of the ad.

preroll

Specifies that the ad is a candidate for preroll insertion. Preroll ad is inserted at the beginning of the playlist.

SEE ALSO

create, delete, edit, list, modify, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-04-10 wam ad-policy(1)

wam application

NAME

application - Configures application for WebAccelerator.

MODULE

wam

SYNTAX

Configure the application component within the wam module using the syntax shown in the following sections.

CREATE/MODIFY

```
create application [name]
modify application [name]
options:
  app-service [[string] | none]
  code [number]
  content-expiration-time [date and time]
  description [string]
  hosts [add | delete | modify | replace-all-with] {
    [ [host name] | [glob] ] {
options:
  app-service [[string] | none]
  code [number]
  subdomain-number-of-http [number]
  subdomain-number-of-https [number]
  subdomain-prefix [string]
  }
}
ibr-adaptive-lifetime [number]
ibr-default-lifetime [number]
ibr-prefix [string]
info-header [none | standard | debug]
multibox [disabled | farm | symmetric]
policy [name]
perf-monitor [enabled | disabled]
perf-monitor-data-retention-period [number]
collect-roi-statistics [ enabled | disabled ]
send-metadata [never | always | uncompressed]
roi-report-email-addresses { string }
roi-report-frequency [ every-month | every-week ]
roi-report-name [string]
roi-report-next-time [date and time]
roi-report-smtp-config [ smtp configuration object name ]
roi-report-collect-statistics { caching-bytes-saved | client-ibred-links |
  compression-bytes-saved | icc-refed-links | inlined-links |
  caching-requests-saved | client-ibred-links-recd | icc-inlined-links |
  image-opt-bytes-saved | minification-bytes-saved }

edit application [ [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

```
reset-stats application
reset-stats application [ [ [name] | [glob] | [regex] ] ... ]
```

DISPLAY

```
list application [name ...]
show running-config application [name ...]
options:
  all-properties
  non-default-properties
  partition
  predefined

show application
show application [name]
options:
  all-properties
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
```

DELETE

```
delete application [name ...]
```

Note: You must remove all references to an application before you can delete it.

DESCRIPTION

You can use the application component to configure the host map, select policies, and set application wide parameters that affect WebAccelerator behavior.

EXAMPLES

```
create application my_app hosts add { host1.com host2.com } policy my_local_policy
```

Creates a WebAccelerator application with a host map consisting of two hosts, host1.com and host2.com, and a local policy set to my_local_policy.

```
modify application my_app remote-policy my_remote_policy
```

Sets my_remote_policy as the remote policy for application my_app.

```
modify application my_app modify hosts { host1.com { subdomain-number-of-http 3 subdomain-prefix abcd } }
```

Sets the number of subdomain hosts to 3 and the subdomain prefix to abcd for host host1.com of WebAccelerator application my_app.

```
modify application my_app roi-report-name "my_report" roi-report-frequency monthly roi-report-next-time now
```

```
roi-report-email-addresses add { someone@domain.com } roi-report-smtp-config smtp-config roi-report-collect-
statistics add { caching_bytes_saved caching_requests_saved }
```

Sets the ROI report name to my_report. The ROI report is set to be sent monthly, and the next time to send the report is set to now. The ROI report will be mailed to someone@domain.com. The SMTP configuration used to send the ROI report will be the predefined configuration of name smtp-config. The ROI report will contain the statistics specified by roi-report-collect-statistics, which in this case would be caching_bytes_saved and caching_requests_saved.

```
delete application my_app
```

Deletes WebAccelerator application my_app.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

code Specifies a numeric non-zero code of the application or application host, which is used for troubleshooting and performance reporting. Each application or application host must have a unique code. If not supplied, it will be generated by the system. Use the keyword generate to specify that the system generate a new unique code.

content-expiration-time

Specifies the date and time that limits how old cached documents can be to still be served from the cache. All documents older than this date and time are considered expired. For example, the following example expires all cached documents of the application my_app:

```
modify application my_app content-expiration-time now
```

description

Specifies the object type description.

hosts

Specifies the list of domain names (host names) that might appear in HTTP requests for your Web application. These are the same host names that DNS has mapped to the server machine on which your WebAccelerator system is running. To map a group or range of requested host names to a single destination host, you can use an asterisk (*) as a wildcard for the first part of the host name.

ibr-adaptive-lifetime

Specifies the adaptive lifetime for Intelligent Browser Referencing in seconds. The default value is 864000 (10 days).

ibr-default-lifetime

Specifies the lifetime for Intelligent Browser Referencing in seconds. The default value is 15724800 (6 months).

ibr-prefix

Specifies a prefix for the Intelligent Browser Referencing tag. The default value is ";wa".

info-header

Enables and controls the appearance of HTTP header X-WA-Info: in responses from WebAccelerator. This header can be used for troubleshooting the WebAccelerator system and for tuning policies. The possible values are:

debug

HTTP header X-WA-Info: is included into responses with standard information, with some additional values to aid WebAccelerator troubleshooting.

none HTTP header X-WA-Info: is not included into responses.

standard

HTTP header X-WA-Info: is included into responses with standard information, such as S-code, policy, and node codes, etc.

multibox

Specifies which type of multibox support is required for this application, if any. Options are disabled, for deployments with an independent WebAccelerator; farm, for farm deployments; and symmetric, for symmetric deployments. When this is not disabled, the application should be shared by a config sync device group containing all devices in the deployment. It enables the broadcast of invalidation messages to other devices in the device group, and, when set to symmetric, also enables symmetric processing of traffic.

partition

Displays the administrative partition within which the application resides.

perf-monitor

Specifies whether performance monitoring for this application is enabled. Enabling performance monitoring on many applications may affect the overall performance of WebAccelerator. The default value is disabled.

perf-monitor-data-retention-period

Specifies the time period in days for how long the performance data must be preserved. The default value is 30 days.

collect-roi-statistics

Specifies whether ROI statistics collection for this application is enabled. The default value is

disabled.

roi-report-name
Specifies the name of ROI statistics report if the statistics collection is enabled and report generation is desired.

roi-report-frequency
Specifies the frequency of ROI statistics report, if the statistics collection is enabled and report generation is desired. The options are every week or month.

roi-report-email-addresses
Specifies the email-addresses to which ROI statistics report will be sent.

roi-report-next-time
Specifies the next time when the ROI statistics report will be sent.

roi-report-smtp-config
Specifies the smtp configuration that will be used to send the scheduled ROI report over email.

roi-report-collect-statistics
Specifies which statistics are to be included in the ROI statistics report.

policy
Specifies the acceleration policy to which you want to assign the new Web application.

predefined
Displays if this application is predefined.

send-metadata
Specifies when Etag HTTP headers will be included into responses. The default value is always.

always
Etag HTTP headers will always be included into responses.

never
Etag HTTP headers will not be included into responses.

uncompressed
Metadata HTTP headers will be included only if response is uncompressed.

subdomain-number-of-http
Specifies the number of HTTP subdomains that you want the WebAccelerator system to generate. The WebAccelerator system uses these additional subdomains only on embedded URLs or links that request images or scripts. The default value is 0.

subdomain-number-of-https
Specifies the number of HTTPS subdomains that you want the WebAccelerator system to generate. The WebAccelerator system uses these additional subdomains only on embedded URLs or links that request images or scripts. The default value is 0.

subdomain-prefix
Specifies the prefix that you want the system to assign to the subdomains. The default value is wa.

For example, if the Requested Host is `www.siterequest.com`, and you select 2 from the HTTP Subdomains box and type `wa` in the Subdomain Prefix box, the WebAccelerator system changes the domain on qualifying embedded URLs and links to use the following domains:

`wa1.www.siterequest.com`
`wa2.www.siterequest.com`

Note: You must configure DNS with these entries, and they must map to the same IP address as the base origin server (`www.siterequest.com` in this example).

SEE ALSO
`create`, `delete`, `edit`, `glob`, `list`, `modify`, `regex`, `reset-stats`, `show`, `tmsh`

COPYRIGHT
No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2014, 2016. All rights reserved.

BIG-IP 2017-05-01 wam application(1)

NAME

domain-list - Configures a list of domains for WebAccelerator for use in inserting DNS prefetch tags.

MODULE

wam

SYNTAX

Configure the domain-list component within the wam resource module using the syntax shown in the following sections.

CREATE/MODIFY

create domain-list [name]

modify domain-list [name]

options:

app-service [[string] | none]

description [[string] | none]

domains

[add | delete | none | replace-all-with] {

[string] ...

}

DISPLAY

list domain-list [name ...]

DELETE

delete domain-list [name ...]

DESCRIPTION

You can use the domain-list component to manage the domain list resources used by the WebAccelerator DNS prefetching feature. A domain-list must be created, then added to the appropriate domain-lists on a WebAccelerator policy node in order for the domains within the domain list to be inserted into a document.

EXAMPLES

```
create domain-list my_domain_list domains add {example.com example2.com}
```

Creates a domain list resource for the domains example.com and example2.com for use in inserting DNS prefetch tags.

```
list domain-list my_domain_list
```

Displays properties of the domain-list resource named my_domain_list.

```
delete domain-list my_domain_list
```

Deletes the domain-list resource named my_domain_list.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

domains

Specifies the domains described by the domain list resource.

SEE ALSO

create, delete, edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2014-06-02 wam domain list(1)

wam object-type

NAME

object-type - Configures object types for WebAccelerator.

MODULE

wam

SYNTAX

Configure the object-type component within the wam module using the syntax shown in the following sections.

CREATE/MODIFY

```
create object-type [name]
modify object-type [name]
options:
  app-service [[string] | none]
  code [ [number] | generate]
  compression [disabled | policy-controlled]
  description [string]
  extensions [add | delete | modify | replace-all-with] {
    [document extension]
    ...
  }
  mime-types [add | delete | modify | replace-all-with] {
    [MIME type]
    ...
  }
  symmetric-compression [ disabled | enabled ]
```

```
edit object-type [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
```

DISPLAY

```
list object-type [ [name] | [glob] | [regex] ] ... ]
show running-config object-type [ [name] | [glob] | [regex] ] ... ]
options:
  all-properties
  non-default-properties
  group
  partition
  predefined
```

DELETE

```
delete object-type [name ...]
```

DESCRIPTION

You can use the object-type component to manage recognized types of objects. These object types are used to classify documents processed by WebAccelerator. A document can be classified by its file extension or MIME type.

EXAMPLES

```
create object-type documents.abcd extensions add { abc abcd } mime-types add { text/abcd text/x-abcd }
description ABCD
```

Creates a WebAccelerator object type named documents.abcd that includes all documents with extensions .abc or .abcd, and MIME types text/abcd or text/x-abcd.

```
delete object-type documents.abcd
```

Deletes the pool named documents.abcd.

```
list object-type documents.abcd
```

Displays properties of the object-type named documents.abcd.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

code Specifies the numeric non-zero code of the object type, which is used troubleshooting and performance reporting. Each object type must have unique code. If not supplied, it will be generated by the system. Use keyword generate to have the system generate a new unique code.

compression

Specifies if this object type supports compression and when it can be enabled. The default value is disabled.

Valid values are:

disabled

Never compresses the response. If you use this option, be aware that it overrides any compression setting configured for the assembly rule that the WebAccelerator system matches to the specified object type. You should use this option only if you want the WebAccelerator system to ignore assembly rules for the specified object type.

policy-controlled

Specifies that compression is controlled by WebAccelerator policy. The compression setting is specified in the assembly rule that the WebAccelerator system matched for this object type. In most cases, you should use this option.

description

Specifies the object type description.

extensions

Specifies the extension the WebAccelerator system should find in the file name or Content-Disposition header of the response, in order to match to the specified object type.

group

Displays the group portion of the name.

mime-types

Specifies the MIME-types that the WebAccelerator system should find in the Content-Type header of the response, in order to match to the specified object type.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify. The name of the object type must be in form group.type where group is used to organize object type based on common usage pattern. for example, documents, binary, pages. The type is used to uniquely identify the object type within a group.

partition

Displays the administrative partition within which the object type resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

symmetric-compression

Specifies whether this object type will be compressed on the WAN link in a symmetric multibox deployment.

SEE ALSO

create, delete, edit, list, modify, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2013. All rights reserved.

BIG-IP 2017-11-17 wam object-type(1)

wam policy

NAME

policy - Configures an acceleration policy for WebAccelerator.

MODULE

wam

SYNTAX

Configure the policy component within the wam module using the syntax shown in the following sections.

CREATE/MODIFY

create policy [name]

modify policy [name]

options:

app-service [[string] | none]

code [integer]

copy-from [name]

description [string]

nodes [add | delete | modify | replace-all-with] {

[name] {

options:

always-proxy [yes | no]

app-service [[string] | none]

assembly-compression [enable | disable]

assembly-compression-ows [enable | disable]

assembly-concatenation [enable | disable]

assembly-concatenation-sets [string] ...

assembly-css-inlining [enable | disable]

assembly-css-inlining-urls [string] ...

assembly-css-reorder [enable | disable]

assembly-css-reorder-cache-size [integer]

assembly-css-reorder-urls [string] ...

assembly-dns-prefetch [enable | disable]

assembly-dns-prefetch-domain-lists [add | delete | replace-all-with] {

[string] ...

```

}
assembly-dns-prefetch-https-enable [enable | disable]
assembly-dns-prefetch-https-automatic [enable | disable]
assembly-ibr [enable | disable]
assembly-image-inlining [enable | disable]
assembly-image-inlining-max-size [integer]
assembly-image-inlining-urls [string ] ...
assembly-js-inlining [enable | disable]
assembly-js-inlining-urls [string ] ...
assembly-js-reorder [enable | disable]
assembly-js-reorder-cache-size [integer]
assembly-js-reorder-urls [string ] ...
assembly-intelligent-client-cache [enable | disable]
assembly-icc-force [enable | disable]
assembly-icc-image-max-size [integer]
assembly-icc-css-inlining-max-size [integer]
assembly-icc-js-inlining-max-size [integer]
assembly-icc-max-num-urls [integer]
assembly-icc-min-client-expiry [integer]
assembly-minification [enable | disable]
assembly-multiconnect [enable | disable]
assembly-on-proxies [enable | disable]
assembly-pdf-linearization [enable | disable]
cache-complete-only [enable | disable]
cache-first-hit [yes | no]
cache-mode [memory-and-disk | memory-only]
cache-priority [low | medium | high]
cache-stand-in-period [integer]
code [integer]
coherency [blade | cluster]
defaults-from [name]
description [string]
jpeg-quality-is-relative [yes | no]
jpeg-quality [integer]
jpeg-strip-keeps-copyright [yes | no]
jpeg-strip-exif [no | yes | if-safe | make-safe]
jpeg-sampling-factor [preserve | 1x1 | 2x1 | 1x2 | 2x2]
jpeg-progressive-encoding [yes | no]
jpegxr-quality [integer]
lifetime-cache-control-extensions
  [add | delete | replace-all-with] {
    [string] ...
  }
lifetime-cache-control-extensions none
lifetime-cache-max-age [integer]
lifetime-honor-ows [yes | no]
lifetime-honor-ows-values
  [add | delete | replace-all-with] {
    [all-values | no-cache | no-store | no-transform |
      max-age | must-revalidate | private | proxy-revalidate |
      s-maxage] ...
  }
lifetime-honor-ows-values none
lifetime-honor-request [yes | no]
lifetime-honor-request-values
  [add | delete | replace-all-with] {
    [all-values | no-cache | no-store | no-transform |
      max-age | max-stale | min-fresh] ...
  }
lifetime-honor-request-values none
lifetime-http-heuristic [percentage]
lifetime-insert-no-cache [yes | no]
lifetime-preserve-response [yes | no]
lifetime-preserve-response-values
  [add | delete | replace-all-with] {
    [all-values | no-cache | no-store | no-transform |
      max-age | must-revalidate | private | proxy-revalidate |
      s-maxage | custom-extension] ...
  }
lifetime-preserve-response-values none
lifetime-response-max-age [integer]
lifetime-response-s-maxage [integer]
lifetime-stand-in-codes
  [add | delete | replace-all-with] {
    [HTTP response code] ...
  }
lifetime-stand-in-codes none
lifetime-use-heuristic [yes | no]
object-max-size [integer | from-profile]
object-min-size [integer | from-profile]
optimize-for-client [yes | no]
options { [hidden | nodelete | nowrite] ...}
order [integer]
response-codes-cached
  [add | delete | replace-all-with] {
    [HTTP response code] ...
  }
}

```

```

viewstate-cache [yes | no]
viewstate-cache-size [integer]
viewstate-tag [string]
video-optimization-fast-start [enable | disable]
video-optimization-max-bitrate [integer]
video-optimization-insert-ad [enable | disable]
video-optimization-preroll-ad [enable | disable]
video-optimization-ad-frequency [integer]
video-acceleration-ad-policy [string]
webp-quality [integer]
matching [add | modify | delete | replace-all-with] {
  [host | path | extension | method:[name] |
  query-param:[name] | unnamed-query-param:[name] |
  path-segment:[name] | cookie:[name] |
  user-agent | referrer | protocol | header:[name] |
  client-ip | content-type] {
    options:
      app-service [[string] | none]
      arg-alias [string]
      arg-direction [left-to-right | right-to-left]
      arg-name [string]
      arg-ordinal [number]
      description [string]
      value-case-sensitive [yes | no]
      values [add | modify | delete | replace-all-with] {
[ [regex] | [string] ] {
      options:
        app-service [[string] | none]
        can-be-empty [yes | no]
        can-be-missing [yes | no]
        invert-match [yes | no]
      }
    }
  }
  values none
}
}
matching none
optimize-image [none | to-jpeg | to-gif | to-png | to-tiff]
png-256-colors [yes | no]
request-queueing [enable | disable]
variation [add | modify | delete | replace-all-with] {
  [host | extension | method:[string] |
  query-param:[name] | unnamed-query-param:[name] |
  path-segment:[name] | cookie:[name] |
  user-agent | referrer | protocol | header:[name] |
  client-ip ] {
    options:
      app-service [[string] | none]
      arg-alias [string]
      arg-all [yes | no]
      arg-ambiguous-as-unnamed [yes | no]
      arg-direction [left-to-right | right-to-left]
      arg-name [string]
      arg-ordinal [number]
      description [string]
      value-case-sensitive [yes | no]
      values [add | modify | delete | replace-all-with] {
[ [regex] | [string] ] {
      options:
        app-service [[string] | none]
        cache-as [same | different]
        can-be-empty [yes | no]
        can-be-missing [yes | no]
        invert-match [yes | no]
        match-all [yes | no]
      }
    }
  }
  values none
}
}
}
variation none
[ proxy | proxy-override ]
[add | modify | delete | replace-all-with] {
  [host | extension | method:[name] |
  query-param:[name] | unnamed-query-param:[name] |
  path-segment:[name] | cookie:[name] |
  user-agent | referrer | protocol | header:[name] |
  client-ip] {
    options:
      app-service [[string] | none]
      arg-alias [string]
      arg-direction [left-to-right | right-to-left]
      arg-name [string]
      arg-ordinal [number]
      description [string]
      value-case-sensitive [yes | no]
      values [add | modify | delete | replace-all-with] {

```

```

    [ [regex] | [string] ] {
options:
app-service [[string] | none]
can-be-empty [yes | no]
can-be-missing [yes | no]
invert-match [yes | no]
}
}
values none
}
}
[ proxy | proxy-override ] none
substitutions [add | modify | delete | replace-all-with] {
[name] {
options:
app-service [[string] | none]
description [string]
dst-alias [string]
dst-direction [left-to-right | right-to-left]
dst-name [string]
dst-ordinal [number]
dst-type [query-param | unnamed-query-param | path-segment]
dst-urls [add | delete | replace-all-with] {
[URI] ...
}
dst-urls none
src-alias [string]
src-direction [left-to-right | right-to-left]
src-name [string]
src-ordinal [number]
src-type
[ randomizer | request-url | query-param |
unnamed-query-param | path-segment ]
src-url [absolute | relative]
}
}
substitutions none
invalidations [add | modify | delete | replace-all-with] {
[name] {
options:
active [yes | no]
app-service [[string] | none]
broadcast [no | yes]
description [string]
cache-content [add | modify | delete | replace-all-with] {
[host | path | extension | method:[name] |
query-param:[name] | unnamed-query-param:[name] |
path-segment:[name] | cookie:[name] |
user-agent | referrer | protocol | header:[name] |
client-ip] {
options:
app-service [[string] | none]
arg-alias [string]
arg-direction [left-to-right | right-to-left]
arg-name [string]
arg-ordinal [number]
description [string]
value-case-sensitive [yes | no]
request-data-alias [string]
request-data-direction [left-to-right | right-to-left]
request-data-name [string]
request-data-ordinal [number]
request-data-type
[ host | path | extension | method |
query-param | unnamed-query-param |
path-segment | cookie | user-agent |
referrer | protocol | header |
client-ip ]
values [add | modify | delete | replace-all-with] {
[ [regex] | [string] ] {
options:
app-service [[string] | none]
can-be-empty [yes | no]
can-be-missing [yes | no]
invert-match [yes | no]
}
}
values none
}
}
}
}
}
partition [name]
publish-build [integer]
publish-comment [string]
published-on [date]

```

Note: Policies can be created only in the Drafts folder. This is required to support publishing functionality. You may create multiple Drafts folders, one for each folder where published policies are going to reside.

DISPLAY

list policy [name ...]
show running-config policy [name ...]
options:
all-properties
non-default-properties
partition
predefined
state

DELETE

delete policy [name ...]

SAVE/LOAD

save policy [name]
load policy [name]
options:
overwrite
file [filename]

PUBLISH

publish policy [name]
options:
publish-comment [string]
publish-build [integer]

Note: Published policies can be deleted, but cannot be modified. The only way to update a published policy is to edit and then publish its development version.

DESCRIPTION

You can use the policy component to manage WebAccelerator acceleration policies. An acceleration policy is a collection of defined rule parameters that dictate how the WebAccelerator system handles HTTP requests and responses. The WebAccelerator system uses two types of rules to manage content: matching rules and acceleration rules. Matching rules are used to classify requests by object type and match the request to a specific acceleration policy. Once matched to an acceleration policy, the WebAccelerator system applies the associated acceleration rules to manage the requests and responses. There are multiple types of acceleration rules: variation, proxy, proxy override, parameter value substitution, and invalidation. The WebAccelerator system ships with several predefined acceleration policies that are optimized for specific Web applications, in addition to several non-application specific policies for general delivery and one for an optional symmetric deployment.

EXAMPLES

Note: For the following examples, the current folder is assumed to be set to /Common.

create policy "Drafts/My Policy"

Creates a new empty policy named My Policy in the folder /Common/Drafts.

create policy "Drafts/My Policy" copy-from "/Common/Generic Policy - Complete"

Creates a new policy My Policy in the folder /Common/Drafts by copying standard system policy /Common/Generic Policy - Complete.

modify policy "Drafts/My Policy" copy-from "/Common/Generic Policy - Complete"

Modifies the policy My Policy by overwriting it with standard system policy /Common/Generic Policy - Complete.

modify policy "Drafts/My Policy" nodes add { "My Node" { default-from Site } }

Adds a new node My Node as the child node of the node Site.

modify policy "Drafts/My Policy" nodes modify { "My Node" { matching add { content-type { values add { pages.other } } } }

Adds a new matching rule into the node My Node. The rule will match content type of the requests to WAM object type pages.other.

publish policy "Drafts/My Policy" publish-comment "Added new node My Node"

Publishes the policy My Policy.

modify policy "Drafts/My Policy" nodes delete { "My Node" }

Deletes the node My Node from the policy My Policy.

delete policy "My Policy"

Deletes the policy My Policy.

save policy "My Policy" file policy.txt

Saves the policy My Policy into the file /var/local/wam/policy.txt.

load policy "Drafts/My Policy" overwrite file /tmp/policy.txt

Loads the policy My Policy from the file /tmp/policy.txt and overwrites the policy if it already exists.

POLICY OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

code Specifies a numeric non-zero code of the policy that is used for troubleshooting and performance reporting. Each policy must have a unique code. If not supplied, it will be generated by the system. Use the keyword generate to specify that the system generate a new unique code.

copy-from

Specifies the name of an existing policy from which to copy all configuration options. If this field is used in the modify command, the configuration options of the existing policy will be replaced with the new ones. The code, state, publish-build, publish-comment, and published-at options are not updated.

description

User defined description of a policy.

nodes

Specifies the collection of policy nodes. Matching rules and acceleration rules for acceleration policies are organized on the Policy Tree, which consists of nodes. The structure of the Policy Tree supports a parent-child relationship. This enables you to easily randomize rules. That is, because a leaf node in a Policy Tree inherits all the rules from its root node and branch node, you can quickly create multiple leaf nodes that contain the same rule parameters by creating a branch with multiple leaf nodes. If you override or create new rules at the branch node level, the WebAccelerator system reproduces those changes to the associated leaf nodes.

partition

Displays the administrative partition within which the policy resides.

publish-build

Specifies the policy build version that was used during policy publishing. If not specified, this number is automatically incremented by the WebAccelerator system.

publish-comment

Specifies the user supplied comment that describes the changes in the policy that is being published.

published-on

Specifies the date and time when this policy was last published.

file Specifies the file name where the policy is going to be saved or loaded from. If a full path is not specified, it is set to /var/local/wam directory.

overwrite

Specifies that the policy file for the save command or the policy component for the load command can be overwritten if it exists.

NODE OPTIONS

always-proxy

Specifies that all requests matching this node must be proxied. If it enabled, proxy rules are not used, even if configured. proxy-override rules still apply.

app-service

Specifies the name of the application service to which this node belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete this node. Only the application service can modify or delete this node.

assembly-compression

Specifies, when enabled, that the WebAccelerator system compresses content for responses, using gzip-encoding. Note that to use this feature, you must set the compress value for the response's object type in the corresponding object-type component, and the client must be able to accept gzip-encoded content. The default value is enabled.

assembly-compression-ows

Specifies, when enabled, that the WebAccelerator system requests gzip-encoded or deflate-encoded content from the origin Web server. Note that the origin Web server will comply only if it supports compression, otherwise it will reply with uncompressed content. The default value is disabled.

assembly-concatenation

Specifies, when enabled, that the WebAccelerator system will perform JavaScript/CSS concatenation in HTML documents. The URLs that may be concatenated are specified using the assembly-concatenation-sets option. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-concatenation-sets

Specifies the concatenation sets that are active for this node. If a URL in the HTML document that belongs to one of the enabled sets is found, it will transformed with concatenation using the URL of the configured set. This is an ordered set, and if the URL exists in multiple active concatenation sets, the first set specified by this option will be used. See the WebAccelerator documentation for more details.

assembly-css-inlining

Specifies, when enabled, that the WebAccelerator system will inline CSS URLs in HTML documents. The CSS URLs that may be inlined are specified using the assembly-css-inlining-urls option. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-css-inlining-urls

Specifies the CSS URLs that may be inlined.

assembly-css-reorder

Specifies, when enabled, that the WebAccelerator system will reorder CSS URLs to the HEAD section of HTML documents. The CSS URLs that may be reordered are specified using the `assembly-css-reorder-urls` option. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-css-reorder-cache-size

Specifies the size of the intermediate cache used to store CSS URLs being reordered. Increasing the size of this cache allows more CSS URLs to be reordered. The default value is 8kB. The maximum value is 8kB.

assembly-css-reorder-urls

Specifies the CSS URLs that may be reordered. The URLs must be fully-qualified and whitespace used to separate URLs. The URLs must correspond to WebAccelerator URL resources created by the command `create wam resource url`. See the help for `wam resource url`.

assembly-dns-prefetch

Specifies, when enabled, that the WebAccelerator system manipulates an HTML document to add DNS prefetch tags at the end of the head. The DNS prefetch tags added are the combined list of domain lists specified in `assembly-dns-prefetch-domain-lists`. DNS prefetch tags will not be inserted in the following conditions: when DNS prefetching is explicitly disabled in the document, either by an HTTP header or by a meta-tag in the head of the HTML document; when the connection is served over HTTPS without `assembly-dns-prefetch-https-enable` enabled; or when the connection is to a client browser that does not support DNS prefetching.

In a document, most browsers will perform DNS prefetching on all domains linked with an HREF. This will speed up performance by having the browser cache possible DNS resolutions before a client clicks on a link, but DNS prefetching cannot automatically occur when a link is created through other means (such as javascript). Inserting DNS prefetch tags addresses this issue.

The default value is disabled.

assembly-dns-prefetch-domain-lists

Specifies the lists of domains that will be inserted into a document. The domain lists must correspond to WebAccelerator domain list resources created by the command `create wam resource domain-list`. See the help for `wam resource domain-list`.

assembly-dns-prefetch-https-enable

Specifies, when enabled, that the WebAccelerator system manipulates an HTML document to add DNS prefetch tags at the end of the head when a document is served over HTTPS when the client browser supports DNS prefetching. By default, most browsers that support DNS prefetching will not do any DNS prefetching on pages served over HTTPS. Enabling `assembly-dns-prefetch-https-enable` will insert a meta-header that will turn on DNS prefetching on the page and a meta-header turning off DNS prefetching for the rest of the page. DNS prefetching cannot be turned on for the rest of an HTML document once the meta-header turning off DNS prefetching is reached.

DNS prefetching is turned off on most browsers serving pages over HTTPS by default as a security measure. DNS prefetching can be used to track which pages are seen over HTTPS by watching the domain resolution requests sent out by the client. According to DNS prefetch standards currently, turning on DNS prefetching on a page will cause all links in the page to have their domains prefetched. This is mitigated by this option with the insertion of an HTTP meta-header turning off DNS prefetching after the DNS tags inserted by the WebAccelerator system. Turning on DNS prefetching for the rest of the page in a request served over HTTPS can be done with the `assembly-dns-prefetch-https-automatic` option.

The default value is disabled.

assembly-dns-prefetch-https-automatic

Specifies, when enabled, that the WebAccelerator system will not insert a meta-tag into an HTML document served over HTTPS turning off DNS prefetching for the rest of a page. By default, most browsers that support DNS prefetching will not do any DNS prefetching on pages served over HTTPS. `assembly-dns-prefetch-https-enable` must be enabled for this option to work.

The default value is disabled.

assembly-ibr

Specifies, when enabled, that the WebAccelerator system manipulates the Web browser cache to reduce requests to your site for relatively static content, such as images and style sheet (CSS) files. The default value is enabled.

assembly-image-inlining

Specifies, when enabled, that the WebAccelerator system will inline image URLs in CSS documents. The image URLs that may be inlined are specified using the `assembly-image-inlining-urls` option. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-image-inlining-max-size

Specifies the maximum size of the image that is allowed to be inlined. The default value is 2kB. The maximum value is 8kB.

assembly-image-inlining-urls

Specifies the image URLs that may be inlined.

assembly-js-inlining

Specifies, when enabled, that the WebAccelerator system will inline JS URLs in HTML documents. The JS URLs that may be inlined are specified using the `assembly-js-inlining-urls` option. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-js-inlining-urls
Specifies the JS URLs that may be inlined.

assembly-js-reorder
Specifies, when enabled, that the WebAccelerator system will reorder JavaScript URLs to the end of HTML documents. The JavaScript URLs that may be reordered are specified using the `assembly-js-reorder-urls` option. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-js-reorder-cache-size
Specifies the size of the intermediate cache used to store JavaScript URLs being reordered. Increasing the size of this cache allows more JavaScript URLs to be reordered. The default value is 8kB. The maximum value is 8kB.

assembly-js-reorder-urls
Specifies the JavaScript URLs that may be reordered. The URLs must be fully-qualified and whitespace used to separate URLs. The URLs must correspond to WebAccelerator URL resources created by the command `create wam resource url`. See the help for `wam resource url`.

assembly-intelligent-client-cache
Specifies, when enabled, that the WebAccelerator system will Intelligent Client Cache HTML documents. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-icc-force
Specifies, when enabled, that the WebAccelerator system will Intelligent Client Cache HTML documents, even if the client does not support HTML5 localstorage. See the WebAccelerator documentation for more details. The default value is disabled.

assembly-icc-image-max-size
Specifies the maximum size of the image that is allowed to be inlined as part of Intelligent Client Caching. The default value is 32kB. The maximum value is 50kB.

assembly-icc-css-max-size
Specifies the maximum size of the CSS that is allowed to be inlined as part of Intelligent Client Caching. The default value is 50kB. The maximum value is 1024kB.

assembly-icc-js-max-size
Specifies the maximum size of the JS that is allowed to be inlined as part of Intelligent Client Caching. The default value is 50kB. The maximum value is 1024kB.

assembly-icc-max-num-urls
Specifies the maximum number of links in an HTML document that are allowed to be inlined as part of Intelligent Client Caching. The default value is 10. The maximum value is 100.

assembly-icc-min-client-expiry
Specifies the minimum client expiry of a resource that is allowed to be inlined as part of Intelligent Client Caching. The default value is 2days.

assembly-minification
Specifies, when enabled, that the WebAccelerator system will minify JavaScript and CSS.

assembly-multiconnect
Specifies, when enabled, that the WebAccelerator system modifies embedded URLs with unique sub-domains that prompt the browser to open more persistent connections for each supported protocol (HTTP or HTTPS). To use this feature, you must configure DNS with the additional domains and map those domains to the same IP address as the base origin server. The default value is enabled.

assembly-on-proxies
Specifies, when enabled, that the WebAccelerator system applies the Content Compression and Intelligent Browser Referencing features (if enabled) to content served to clients, even if the content is not served from the WebAccelerator system's cache. Enable this option if you are using the Content Compression or Intelligent Browser Referencing features. The default value is enabled.

assembly-pdf-linearization
Specifies, when enabled, that the WebAccelerator system applies linearization on PDF documents, if the documents match the node matching rules. PDF linearization transforms the document to include the index of the pages in the beginning. This allows Web browsers to load and show specific pages rather than a whole document. See the WebAccelerator documentation for more details. The default value is disabled.

optimize-image
Specifies whether image optimization should be applied and the format conversion to use. Each of the 4 supported formats (JPEG, PNG, GIF, TIFF) can be converted to any of the others. Images using a capability unique to one format may lose that feature when converted to a format that does not support it. (For example, animated GIFs or multipage-TIFFs will have only the first image when converted to PNG or JPEG). Transparency will be lost when converting from GIF or PNG to JPEG. TIFF is a container for many different image formats so the results will be best-effort and may not list completely.

A converted image will likely have a different number of bytes after conversion. Some conversions are likely to produce fewer bytes; however, a requested conversion will be done even if it results in more bytes (for consistency). For example, you may want to offer multiple formats of an image without storing them all on a server.

A correct Content-Type header will be generated for converted images, but HTML files will not be rewritten.

optimize-for-client Specifies whether to allow conversion to a format and/or size which is optimum for the specific client making the request but which, if saved by that client and later sent elsewhere, might not be

appropriate.

webp-quality WebP is a "lossy" compression format. This means when you convert an image to a WebP and then convert it back, you will not get back exactly the same image you started with. Compression changes the amount of information stored (and therefore the number of bytes), but not the image dimensions (the number of pixels). The **webp-quality** attribute represents the absolute quality of the WebP produced. Compression (quality) is represented as a number between 1-100 where 1 is minimal quality, but small, and 100 is high-quality, but large. For most images, useful values of quality will be from about 30-70.

jpegxr-quality JPEG-XR is a "lossy" compression format. This means when you convert an image to a JPEG-XR and then convert it back, you will not get back exactly the same image you started with. Compression changes the amount of information stored (and therefore the number of bytes), but not the image dimensions (the number of pixels). The **jpegxr-quality** attribute represents the absolute quality of the JPEG-XR produced. Compression (quality) is represented as a number between 1-100 where 1 is minimal quality, but small, and 100 is high-quality, but large. For most images, useful values of quality will be from about 5-30.

jpeg-quality-is-relative =item **jpeg-quality**

JPEG is a "lossy" compression format. This means when you convert an image to a JPEG and then convert it back, you will not get back exactly the same image you started with. Compression changes the amount of information stored (and therefore the number of bytes), but not the image dimensions (the number of pixels). When **jpeg-quality-is-relative** is set to no, the **jpeg-quality** attribute represents the absolute quality of the JPEG produced. Compression (quality) is represented as a number between 1-100 where 1 is minimal quality, but small, and 100 is high-quality, but large. For most images, useful values of quality will be from about 30-100. Because information once lost cannot be regained, converting a low-quality JPEG to a higher quality is pointless and image optimization will prevent that (by not changing the original to a higher JPEG quality).

You might be unable to choose a specific absolute quality for JPEG images. When **jpeg-quality-is-relative** is set to yes, the relative JPEG quality setting is enabled. In this case, **jpeg-quality** is a percentage (a number between 1-100) that when multiplied by each JPEG's original quality, becomes its optimized quality.

jpeg-strip-exif

JPEG files have a header (called EXIF) that contains optional data such as a date, time, camera model, exposure settings, and so on. The EXIF header can also contain a color profile, which is required when included. EXIF headers can be small or large. Unless they contain a color profile, they do not affect displaying the image, and so can be removed if the loss of the information they contain is acceptable. There are four options for this setting:

no Leaves any EXIF headers alone.

yes Always strips EXIF headers.

if-safe

Only strips EXIF headers if they do not have color profiles (ensures that images display properly).

make-safe

Applies the color profile and then strips the EXIF header (typically decreases image file size). Applying a color profile requires additional CPU time.

jpeg-strip-keeps-copyright This setting affects the meaning of **jpeg-strip-exif**. If it is set, stripping the EXIF header will strip everything except the Copyright notice (if one is present).

jpeg-sampling-factor

Sets the sampling factor to be used when producing JPEG images. The default value is **preserve**, which matches the original file. You can also explicitly specify this option, as it can sometimes improve compression.

jpeg-progressive-encoding

When enabled, progressive encoding will be used in JPEG images. For large JPEG files, this can improve compression. When this is enabled, it will be applied only if the file is large enough to improve compression.

png-256-colors

It is often possible to significantly reduce the size of PNG files without changing their appearance very much by reducing the number of colors to 256 optimally selected values. This optimization is enabled when **png-256-colors** is set to **yes**.

cache-complete-only

Specifies, when enabled, that the WebAccelerator system caches HTML pages only if the HTML code within the page contains begin and end tags. When disabled, the WebAccelerator system reviews HTTP response headers to determine if the information contained on the page is complete. The default value is **enabled**.

cache-first-hit

Specifies that the first response should be cached according to the policy caching settings. When this is off, the response is cached when more than one request for the document has been seen. Turning this on can reduce cache churn for unpopular documents. The default value is **no**.

cache-mode

Specifies how where the cached documents will be stored. The default value is **memory-and-disk**. Possible values are:

memory-and-disk

The cached documents will be stored in memory or on disk.

memory-only

The cached documents will be stored in memory only.

cache-priority

Specifies the cache admission priority of documents matching the policy node. Documents with high priority are more likely to be admitted into the cache. The default value is **medium**. Possible values are:

low Documents will have low priority.

medium

Documents will have medium priority.

high Documents will have high priority.

cache-stand-in-period

Specifies the amount of time that the WebAccelerator system continues to serve content from cache if the origin Web server does not respond to the WebAccelerator system's requests for fresh content. The default value is 0 (zero), which means the WebAccelerator system responds to requests for expired content with a HTTP 404 error.

code Specifies a numeric non-zero code for the node that is used for troubleshooting and performance reporting. All nodes must have unique codes within the policy. If not supplied, the code will be generated by the system. Use the keyword generate to specify that the system generate a new unique code.

coherency

Specifies if the WebAccelerator system will attempt to keep content matching the associated node in sync across the blades of a cluster. The default behavior is to keep content in sync.

blade

The cached documents will not be kept coherent across blades. This causes each blade to have its own copy of a given cached document.

cluster

The cached documents will be kept coherent across blades. This causes the cluster to have single version of a given cached document.

defaults-from

Specifies the node that you want to use as the parent node. Your new node inherits all options and values from the parent node specified. The default value is none, which means this is a root node.

description

User defined description of a node.

invalidations

Specifies the collection of invalidations rules. Invalidations rules enable you to expire cached content before it has reached its time-to-live (TTL) value. This is useful when content updates are event-driven, such as when an item is added to a shopping cart, a request contains a new auction bid, or a poster has submitted content on a forum thread. Invalidations rules can be created only on leaf nodes.

lifetime-cache-control-extensions

Enables you to configure extension tokens to be added to the cache-control header of HTTP response. The WebAccelerator system does not process any of these extensions. It is possible that the origin Web server will send cache-control extensions as well. You can choose whether to preserve them by including the custom-extension in the lifetime-preserve-response-values list.

lifetime-cache-max-age

Specifies the amount of time that the WebAccelerator system serves content from the cache before requesting fresh content from the origin Web server. The default value is 4 hours.

lifetime-honor-ows

Specifies, if enabled, that the WebAccelerator system honors certain cache-control directives from the origin Web server response to determine cache lifetime. The default value is disabled.

lifetime-honor-ows-values

Specifies which Cache-Control directive from the origin Web server response will determine cache lifetime. Available directives are all-values, private, no-cache, no-store, must-revalidate, proxy-revalidate, max-age, s-maxage, and expires. This option is only effective if lifetime-honor-ows is enabled.

lifetime-honor-request

Specifies, if enabled, that the WebAccelerator system honors certain Cache-Control directives from the client's browser request to determine cache lifetime. The default value is enabled.

lifetime-honor-request-values

Specifies which cache-control directive from client's browser request will determine cache lifetime. Available directives are all-values, no-cache, no-store, max-age, max-stale, and min-fresh. This option is only effective if lifetime-honor-request is enabled. The default values are max-age, max-stale, and min-fresh.

lifetime-http-heuristic

Specifies the percentage, based on the HTTP Last-Modified header, that the WebAccelerator system uses to compute TTL values for cached content. For example, if content was modified 30 days ago and the lifetime-http-heuristic option is set to 50%, the WebAccelerator system caches the content for 15 days. This option is applicable only if you use the HTTP Last-Modified headers to identify content lifetime. The default value is 50%. This option is effective only if lifetime-use-heuristic is enabled.

lifetime-insert-no-cache

Specifies, when enabled, that the WebAccelerator system inserts a no-cache directive into the HTTP Cache-Control header, which stops the client's browser from locally caching content. This value overrides the HTTP Cache-Control header cache directives sent to the client by the origin Web server.

lifetime-preserve-response

Specifies, if enabled, that the WebAccelerator system preserves certain Cache-Control directives from the

origin Web server and includes them into client's browser response. The default value is enabled.

lifetime-preserve-response-values

Specifies which Cache-Control directive from the origin web server response will be preserved in response to the client's web browser. Available directives are all-values, private, no-cache, no-store, must-revalidate, proxy-revalidate, max-age, s-maxage, expires, and custom-extension. This option is only effective if lifetime-preserve-response is enabled. The default value is all-values.

lifetime-response-max-age

Specifies, when enabled, the amount of time that the client's browser should locally store content. This value overrides the max-age and expires the directives in the HTTP Cache-Control header that are sent to the client by the origin web server, only if the new value for the max-age is greater than the value supplied by the origin web server. Modify this value only if there is an acceptable trade off between the freshness of the content served to clients and overall site performance.

lifetime-response-s-maxage

Specifies, when enabled, the amount of time that the client's browser should locally store shared content. This value overrides the s-maxage and expires the directives in the HTTP Cache-Control header that are sent to the client by the origin web server, only if the new value for the s-maxage is greater than the value supplied by the origin web server. Modify this value only if there is an acceptable trade off between the freshness of the shared content served to clients and overall site performance.

lifetime-stand-in-codes

Specifies that the WebAccelerator system is allowed to serve stale content from the cache if it is not able to re-validate its freshness with the origin web server. The WebAccelerator system serves invalid content to the downstream proxies or clients if the response code from the origin web server matches one of codes specified with this option. This option is effective only if cache-stand-in-period has a non-zero value. The default values are 404, 500, and 504.

lifetime-use-heuristic

Specifies, when enabled, that the WebAccelerator system uses the percentage from lifetime-use-heuristic option to compute TTL values for cached content. The default value is no.

matching

Specifies the collection of matching rules. The rules consist of the HTTP request data type parameters that the WebAccelerator system uses to match an incoming HTTP request to a specified node. The following types of HTTP parameters are available for matching rules: host, path, extension, query-param, unnamed-query-param, path-segment, cookie, user-agent, referrer, protocol, method, header, client-ip, and content-type.

object-min-size

Specifies the minimum object size required in order for content matching the associated node to be eligible for caching. The default behavior is to use the minimum object size specified by the associated web-acceleration profile.

object-max-size

Specifies the maximum object size allowed for content matching the associated node in order to be eligible for caching. The default behavior is to use the maximum object size specified by the associated web-acceleration profile.

order

Specifies the order of the node in the Policy Tree. All nodes in the policy must have an order. The order numbers are sequential, starting from 2. Orders 0 and 1 are reserved for internal use. The child node orders must be greater than the order of their parent node. You can change the order of the nodes by updating the order option of the node that you would like to move. The system honors the specified order if it falls within the range of sibling node orders. Otherwise, the system picks the closest valid order number. The remaining nodes are automatically re-ordered to free requested order number. The node order is also used as a last resort to determine which node to use when multiple nodes match the request. The node with a lower order wins. New nodes have their order assigned automatically to make them last among their siblings.

proxy

Specifies the collection of proxy rules. In general, proxy rules options are relevant to only requests that match their node, rather than to matched responses. The following types of HTTP parameters are available for proxy rules: host, query-param, unnamed-query-param, path-segment, cookie, user-agent, referrer, protocol, method, header, and client-ip.

proxy-override

Specifies the collection of proxy override rules. You can define proxy override rules and associated conditions under which the WebAccelerator system should ignore proxying rules options. The following types of HTTP parameters are available for proxy override rules: host, query-param, unnamed-query-param, path-segment, cookie, user-agent, referrer, protocol, method, header, and client-ip.

request-queueing

Specifies, when enabled, that the WebAccelerator system will queue requests for expired or new documents and proxy fewer requests to the origin web server (OWS). If the response is cachable, the response will be served to all waiting requests; if not, the waiting requests will proxy normally.

response-codes-cached

Specifies the collection of HTTP response codes that determine whether the WebAccelerator system should cache the content. The valid codes are 300, 301, 302, 307, and 410. The codes 200, 201, 203, and 207 are included into the list implicitly. The default values are 300 and 301.

substitutions

Specifies the collection of parameter value substitution rules. Some requested pages include hyperlinks that require that specific information appear in the response. You can configure parameter value substitution so that when a query parameter contains identification information for a sites visitors, it

prompts the WebAccelerator system to serve different content for the request, based on the specific visitor. Conversely, if parameter value substitution is not configured, the WebAccelerator system uses the value that it cached for the original request, for all subsequent requests after the first, even if the subsequent requests have different values that should be used in the response.

If you configure parameter value substitution, the WebAccelerator system changes the targeted parameters value on the page served from the cache, so that the parameter you specify appears on the URL embedded in that page.

variation

Specifies the collection of variation rules. When the WebAccelerator system caches responses from the origin web server, it uses certain HTTP request parameters to create a Unique Content Identifier (UCI). The WebAccelerator system stores the UCI in the form of a compiled response and uses the UCI to easily match future requests to the correct content in its cache. You can configure variation rules to add or modify the parameters on which the WebAccelerator system bases its caching process. If the WebAccelerator system receives two requests that are identical except for the value of a query parameter defined in the variation rule, it creates a different UCI for each, and caches each response under its unique UCI. The following types of HTTP parameters are available for variation rules: host, query-param, unnamed-query-param, path-segment, cookie, user-agent, referrer, protocol, method, header, and client-ip.

viewstate-cache

Specifies, when enabled, that the WebAccelerator system accelerates requests and responses for Web form objects that are generated by ASP.NET web applications. Because the file size of forms can be significant, the WebAccelerator system is able to cache and substitute values, thus reducing the file size and achieving faster performance.

viewstate-cache-size

Specifies the size of the ViewState object cache in kilobytes. The default value is 100 kilobytes.

viewstate-tag

Specifies the name of the web form field where the ViewState object is stored. The default value is `__VIEWSTATE`.

video-optimization-fast-start

Specifies when enabled, that the WebAccelerator system optimizes video by prefetching.

video-optimization-max-bitrate

Specifies, the maximum bitrate of video that can be allowed in kilobits per sec. The default value is 0.

video-optimization-insert-ad

Specifies, when enabled, that the WebAccelerator system can insert ad into the video.

video-optimization-preroll-ad

Specifies, when enabled, that the WebAccelerator system can insert ad at the beginning of the video.

video-optimization-ad-frequency

Specifies the frequency of ad insertion. Units in seconds.

video-optimization-ad-policy

Specifies the ad policy applicable when processing the video.

type Displays the node type. The possible types are:

branch

The branch nodes exist only for the purpose of propagating rule parameters to leaf nodes. The WebAccelerator system does not perform matching against branch nodes. Branch nodes can have multiple leaf (child) nodes, as well as child branch nodes.

A leaf node inherits rule parameters from its parent branch node. The WebAccelerator system performs matching only against leaf nodes, and then applies the leaf nodes corresponding acceleration rules to the request.

HTTP PARAMETERS

Both matching and acceleration rules are identified by the type, and optionally, by the name of HTTP parameters that are used inside the rules. The following types of HTTP parameters are available:

content-type

A rule that uses the content-type parameter is based on type definitions in the object-type components. Unlike the HTTP request data types, a matching rule based on content type is specific to the content type parameter that the WebAccelerator system generates for a response. You specify the regular expression that you want a response's content type to match.

client-ip

A rule that uses the client IP parameter is based on the IP address of the client making the request. The IP address, however, may not always be the address of the client that originated the request. For example, if the client goes through a proxy server, the IP address is the IP address of the proxy server, rather than the client IP address that originated the request. If several clients use a specific proxy server, they all appear to come from the same IP address.

cookie:[name]

A rule that uses the cookie parameter is based on a particular cookie that you identify by name, and for which you provide a value to match against. This value is usually literal and must appear on the cookie in the request or in a regular expression that matches the request's cookie that appears on the cookie HTTP request headers. These are the same names you use to set the cookies, using the HTTP Set-Cookie response headers. The HTTP request can contain multiple cookies, and the rule identifier must include the name of the cookie separated with colon (:).

extension

A rule that uses the extension parameter is based on the value that follows the far-right period, in the far-right segment key of the URL path.

header:[name]

A rule that uses the header parameter is based on a particular header that you identify by name and for which you provide a value to match against. You can use an HTTP request data type header parameter to create rules based on any request header other than one of the recognized HTTP request data types. The HTTP request can contain multiple headers, and the rule identifier must include the name of the header separated with colon (:).

host A rule that uses the host parameter is based on the value provided for the HTTP Host request header field. This header field describes the DNS name that the HTTP request is using.

method

A rule that uses the method parameter is based on whether the request uses the GET or POST method.

query-param:[name]

A rule that uses the query parameter is based on a particular query parameter that you identify by name and for which you provide a value to match against. The value is usually literal and must appear on the query parameter in the request, or in a regular expression that matches the requests query parameter value. The query parameter can be in a request that uses GET or POST methods. The HTTP request can contain multiple query parameters, and the rule identifier must include the name of the header separated with colon (:).

path A rule that uses the path parameter is based on the path portion of the URI. The path is defined as everything in the URL after the host and up to the end of the URL, or up to the question mark (whichever comes first).

path-segment:[name]

A segment is the portion of a URI path that is delimited by a forward slash (/). For example, in the path: /apps/search/full/complex.jsp, apps, search, full, and complex.jsp all represent path segments. The path can contain multiple segments so the rule identifier must include the name of the segment separated with colon (:). The name can be a segment ordinal or some other string to distinguish it from other segments rules in the same node.

protocol

A rule that uses the protocol parameter is based on whether the request uses the HTTP or HTTPS protocol.

referrer

A rule that uses the referrer parameter is based on the value provided for the HTTP Referer in the request header. (Note the misspelling of Referer. This spelling is defined for this request header in all versions of the HTTP specification.) This header provides the URL location that referred the client to the page that the client is requesting. You do not typically base rules on the Referer request header, unless you want your sites behavior to be dependent on the specific referrer. For example, one implementation would be for sites that provide different branding for their pages based on the user's web portal or search engine.

unnamed-query-param:[name]

An unnamed query parameter is a query parameter that has no equal sign. That is, only the query parameter value is provided in the URL of the request. The HTTP request may contain multiple unnamed query parameters so the rule identifier must include the name of it separated with colon (:). The name can be the ordinal of unnamed query parameter or some other string that can make it distinguishable from other unnamed query parameter rules in the same node.

user-agent

A rule that uses the user agent parameter is based on the value provided for the HTTP User-Agent in the request header, which identifies the browser that sent the request.

RULE OPTIONS

active

Specifies, when enabled, that the invalidation trigger rule is enabled. You can use this option to disable a specific invalidation trigger rule temporary, without removing it from the policy.

arg-all

Specifies, when enabled, that the rule matches all HTTP parameters of this type rather than one identified by arg-name or arg-ordinal. This option is applicable to variation rules query-param, unnamed-query-param, path-segment, cookie, and header. Such rules serve as a fallback case for defining document variation. All root nodes must include one variation rule of each type with this option enabled. The default value is disabled.

arg-alias

src-alias

dst-alias

request-data-alias

Specifies the user supplied alias for rules that use ordinals to identify HTTP request data. These include the unnamed-query-param and path-segment rules. The src-alias and dst-alias options are used in parameter value substitution rules to define aliases for the source and target definitions correspondingly. The request-data-alias option defines an alias for the invalidation trigger rules.

arg-direction

src-direction

dst-direction

request-data-direction

Specifies the direction that the WebAccelerator system uses to count the ordinal of path-segment. The src-direction and dst-direction options are used in parameter value substitution rules to define the ordinal direction for the source and target definitions correspondingly. The request-data-direction

option defines the ordinal direction for the invalidation trigger rules. The default value is left-to-right. The possible values are:

left-to-right

The path segment is counted form left to right.

right-to-left

The path segment is counted form right to left.

arg-name
src-name
dst-name
request-data-name

Specifies the name of the parameter type for query-param, cookie, and header. If not specified, arg-name option is initialized from the rule name. This option is not effective if arg-all is enabled. The src-name and dst-dst options are used in parameter value substitution rules to define the parameter name for the source and target definitions correspondingly. The request-data-name option defines the parameter name for the invalidation trigger rules.

arg-ordinal
src-ordinal
dst-ordinal
request-data-ordinal

Specifies, in the form of a number, the location of a parameter for unnamed-query-param and path-segment rules. The numbering starts at 1 and follows the direction specified in the corresponding direction option. This option is not effective if arg-all is enabled. The src-ordinal and dst-ordinal options are used in parameter value substitution rules to define the parameter ordinal for the source and target definitions correspondingly. The request-data-ordinal option defines the parameter ordinal for the invalidation trigger rules.

broadcast

Specifies whether a triggered invalidation rule is broadcast to other members of a multibox deployment. This option is only effective when the application using this policy has multibox set to farm or symmetric.

cache-content

Specifies the parameter for which the WebAccelerator system must obtain fresh content when the invalidations rule is triggered. The available request types are: host, path, extension, query-param, unnamed-query-param, path-segment, cookie, user-agent, referrer, protocol, method, header, and client-ip.

Note: You must select and configure the path parameter for the cached content to invalidate, or the invalidations rule will fail to trigger. All other parameters are optional.

description

User-defined description of a rule.

dst-type

Specifies the HTTP parameter type to use as target definition for the request value substitution rule. A target definition contains a value in the embedded URL that you want the WebAccelerator system to replace with the value that you specified for the source definition, during assembly. The possible values are:

path-segment

Specifies that the WebAccelerator system targets the URL parameter, as specified by the dst-ordinal and dst-direction you define.

query-param

Specifies that the WebAccelerator system targets the URL parameter, as specified by the dst-name you define.

unnamed-query-param

Specifies that the WebAccelerator system substitutes the URL parameter, as specified by the dst-ordinal you define.

dst-urls

Specifies the collection of URLs in the request for which you want the WebAccelerator system to replace content.

request

Specifies a parameter in the request that triggers the invalidations rule. The available request types are: host, path, extension, query-param, unnamed-query-param, path-segment, cookie, user-agent, referrer, protocol, method, header, and client-ip.

Note: You must select and configure the path parameter for the request header criteria, or the invalidations rule will fail to trigger. All other parameters are optional.

request-data-type

Specifies the HTTP request parameter value that the WebAccelerator system should find in its cache and for which it should request updated content from the origin Web server. The default value is undefined.

The following types of HTTP parameters are available:

host
query-param
unnamed-query-param
path-segment
cookie
user-agent

**referrer
header
client-ip**

Specifies that the WebAccelerator system should use the corresponding value from the request that triggered the invalidation. Additional data, if required to identify the value, must be specified in the request-data-name, request-data-ordinal, and request-data-direction options. The values option is ignored.

undefined

Specifies that the WebAccelerator system should not use any values from the request that triggered the invalidation. You must add a value into the values option with which to compare the cached content.

src-type

Specifies the HTTP parameter type to use as source definition for the request value substitution rule. A source definition contains the value that the WebAccelerator system embeds in the URL, in place of the cached (target definition) value, during substitution. Typically, the source definition is a specific request element, such as a particular query parameter; however, you can specify another source type, such as a random number. The possible values are:

path-segment

Specifies that the WebAccelerator system substitutes the URL parameter, as specified by the src-ordinal and src-direction options you define.

query-param

Specifies that the WebAccelerator system substitutes the URL parameter, as specified by the src-name option you define.

randomizer

Specifies that the WebAccelerator system generates a random number and places that number on the targeted location in an embedded URL.

request-url

Specifies that the WebAccelerator system is limited to target-specific URLs embedded in a page, as defined in the prefix that an embedded URL must match before the WebAccelerator system performs substitution. If you use the request URL as the source, the WebAccelerator system uses the entire request URL as the value to substitute.

unnamed-query-param

Specifies that the WebAccelerator system substitutes the URL parameter, as specified by the src-ordinal option you define.

src-url

Specifies whether the request URL is a relative URL or an absolute URL. The default value is absolute.

value-case-sensitive

Specifies, when enabled, that the HTTP parameter must be matched against supplied value(s) in case sensitive manner. The default value is no.

values

Values are a collection of rule parameters that enable you to specify different parameter values for the same rule. Most rules allow only one value, while variation rules support multiple values. Each value can prompt a different behavior by the WebAccelerator system. All variation rules must include at least one value with match-all option enabled. A value can be represented by actual string, regex, or multiple strings, or regexes separated by space ().

RULE VALUE OPTIONS

can-be-empty

Specifies, when enabled, that the defined HTTP request parameter is included in the request, but has no value (is an empty string). The default value is no.

can-be-missing

Specifies, when enabled, that the defined HTTP request parameter is absent from the request. The default value is no.

invert-match

Specifies, when enabled, that the defined HTTP request parameter does not match the associated regular expression that you defined. The default value is no.

match-all

Specifies, when enabled, that the defined HTTP request parameter matches all possible values. This option is available only for variation rule values as a fallback case. Each variation rule must have at least one value with this option enabled. The default value is no.

cache-as

Specifies whether the associated value should prompt the WebAccelerator system to reply to matched requests with the same or different content. This option is available only for variation rule values.

SEE ALSO

create, delete, edit, list, modify, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2013, 2016. All rights reserved.

wam resource concat-set

NAME

concat-set - Configures concatenation sets for WebAccelerator for use in JavaScript/CSS concatenation

MODULE

wam resource

SYNTAX

Configure the concat-set within the wam resource module using the syntax shown in the following sections.

CREATE/MODIFY

```
create concat-set [name]
modify concat-set [name]
options:
  app-service [[string] | none]
  url [url]
  type [css|js]
  members [string ] ...
```

DISPLAY

```
list concat-set [name ...]
```

DELETE

```
delete concat-set [name ...]
```

DESCRIPTION

You can use the concat-set component to manage the concatenation sets used by the WebAccelerator JavaScript and CSS concatenation feature. A concatenation set must be created, then enabled and activated in the configuration on a WebAccelerator policy node.

EXAMPLES

```
create concat-set testSet url http://www.example.com/concatSet.css type css
```

Creates a set whose URL will be http://www.example.com/concatSet.css for use in concatenation.

```
list concat-set testSet
```

Displays properties of the concatenation set named testSet.

```
delete concat-set testSet
```

Deletes concatenation set named testSet.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

url Specifies the URL that will be used to generate the concatenated link.

type Either "css" or "js". Specifies whether the set is to be used for CSS or JavaScript concatenation.

members

Specifies the members of this set. The set members are the URL resources that are defined by the wam resource url.

SEE ALSO

create, delete, edit, list, modify, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

wam resource domain-list

NAME

domain-list - Configures a list of domains for WebAccelerator for use in inserting DNS prefetch tags.

MODULE

wam resource

SYNTAX

Configure the domain-list component within the wam resource module using the syntax shown in the following sections.

CREATE/MODIFY

create domain-list [name]

modify domain-list [name]

options:

app-service [[string] | none]

description [[string] | none]

domains

[add | delete | none | replace-all-with] {

[string] ...

}

DISPLAY

list domain-list [name ...]

DELETE

delete domain-list [name ...]

DESCRIPTION

You can use the domain-list component to manage the domain list resources used by the WebAccelerator DNS prefetching feature. A domain-list must be created, then added to the appropriate domain-lists on a WebAccelerator policy node in order for the domains within the domain list to be inserted into a document.

EXAMPLES

```
create domain-list my_domain_list domains add {example.com example2.com}
```

Creates a domain list resource for the domains example.com and example2.com for use in inserting DNS prefetch tags.

```
list domain-list my_domain_list
```

Displays properties of the domain-list resource named my_domain_list.

```
delete domain-list my_domain_list
```

Deletes the domain-list resource named my_domain_list.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

domains

Specifies the domains described by the domain list resource.

SEE ALSO

create, delete, edit, list, modify, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-12-23 wam resource domain-list(1)

wam resource url

NAME

url - Configures a URL resource for WebAccelerator for use in reordering whitelists

MODULE

wam resource

SYNTAX

Configure the url component within the wam resource module using the syntax shown in the following sections.

CREATE/MODIFY

create url [name]

modify url [name]

options:

app-service [[string] | none]

url [url]

type [css|js]

DISPLAY

list url [name ...]

DELETE

delete url [name ...]

DESCRIPTION

You can use the url component to manage the URL resources used by the WebAccelerator JavaScript and CSS reordering features. A URL resource must be created, then added to the appropriate whitelist on a WebAccelerator policy node in order for the corresponding URL to be reordered.

EXAMPLES

```
create url test.css url http://www.example.com/test.css type css
```

Creates a URL resource for the URL <http://www.example.com/test.css> for use in CSS reordering whitelists.

```
list url test.css
```

Displays properties of the URL resource named test.css.

```
delete url test.css
```

Deletes the URL resource named test.css.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

url Specifies the URL described by the URL resource.

type Either "css" or "js". Specifies whether the URL resource is to be used for CSS or JavaScript reordering.

SEE ALSO

create, delete, edit, list, modify, show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2013. All rights reserved.

BIG-IP 2013-04-12 wam resource url(1)

wam roi-statistics

NAME

roi-statistics - Provides ROI statistics for WAM application.

MODULE

wam

SYNTAX

Provides ROI statistics for configured application within the wam module.

DISPLAY

```
show roi-statistics
show roi-statistics [application-name]
options:
  (default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)
  field-fmt
```

DESCRIPTION

You can use the roi-statistics component to view the ROI statistics of configured WAM application.

EXAMPLES

```
show roi-statistics my_application
```

The show command will display ROI statistics for the configured WAM application.

SEE ALSO

show, tmsb

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2011-2014. All rights reserved.

BIG-IP 2014-07-02 wam roi-statistics(1)

wom

wom advertised-route

NAME

advertised-route - Configures a route advertised by the local endpoint to remote endpoints for WAN optimization.

MODULE

wom

SYNTAX

Configure the advertised-route component within the wom module using the syntax in the following sections.

CREATE/MODIFY

```
create advertised-route [name]
modify advertised-route [name | all]
options:
  app-service [[string] | none]
  description [string]
  dest [ip address/netmask]
  include [disabled | enabled]
  label [value]
  metric [integer]
  origin [configured | discovered | manually-saved | persistable]
```

DISPLAY

```
list advertised-route
show advertised-route
options:
  running-config
```

DELETE

```
delete advertised-route [name]
```

DESCRIPTION

You can use the advertised-route component to configure a subnet that the system can reach through the local endpoint. You can specify a netmask or use slash format.

Routes are advertised to all connected WAN Optimization Managers. The remote endpoints use the subnet configuration information to determine peer routing and optimization actions.

EXAMPLES

```
list advertised-route all
```

Displays all endpoint advertised routes for the local WAN Optimization Manager.

`delete advertised-route adv_rt2`

Deletes the advertised route `adv_rt2`.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the `strict-updates` option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

dest Specifies the IP address and netmask of the advertised route.

include

Enables or disables the inclusion of this route in the optimization of traffic. This option allows you to define a subset of IP addresses to exclude from optimization within a larger included subnet. An excluded endpoint advertised route must be a valid address range subset of an included endpoint advertised route. The default is enabled.

label

Specifies an optional descriptive label for this route.

metric

Specifies a routing number to select between WAN Optimization Manager pairs. The higher the number, the more expensive the route in terms of resources. Not currently implemented.

origin

Specifies whether the route was discovered automatically or configured manually. You can change the origin from `discovered` to `persistable`, if you want to save the route to the file `bigip_local.conf` when you use the command `save config`. After you run the command `save config`, this attribute changes to `manually saved`. Endpoints that have the attribute `discovered` are not saved to the file `bigip_local.conf`.

The options are:

configured

Indicates that you manually configured this route. The system automatically sets this value, and you cannot change it.

discovered

Indicates that the system automatically discovered this route. Note that route for which the value of the origin property is `discovered` are not saved to the file `bigip_local.conf`.

manually-saved

After you run the command `save / sys config`, the value of the origin property that was set to `persistable` changes to `manually-saved`. Note that after the system changes the value to `manually-saved`, you cannot change it again.

persistable

Change the origin from `discovered` to `persistable`, if you want to save the route to the file `bigip_local.conf` when you use the command `save / sys config`.

SEE ALSO

`create`, `delete`, `list`, `wom local-endpoint`, `modify`, `wom remote-endpoint`, `wom server-discovery`, `show`, `tmsh`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2017-03-31 wom advertised-route(1)

wom deduplication

NAME

`deduplication` - Configures symmetric data deduplication for WAN optimization.

MODULE

`wom`

SYNTAX

Configure the deduplication component within the wom module using the syntax in the following sections.

MODIFY

```
modify deduplication
options:
  codec [sdd-v2 | sdd-v3]
  [disabled | enabled]
  max-endpoint-count [integer]
```

DISPLAY

```
list deduplication
show running-config deduplication
options:
  dictionary-size
  one-line
```

DESCRIPTION

You can use the deduplication component to configure symmetric data deduplication, which compresses data over the WAN by identifying and removing repetitive data patterns.

EXAMPLES

```
list deduplication
```

Displays all the deduplication settings.

```
modify deduplication max-endpoint-count 4
```

Sets the maximum number of remote endpoints to 4.

OPTIONS

codec
Specifies which algorithm the system uses for deduplication.

The options are:

sdd-v2

Used for low number of spokes, such as for data center to data center replication.

sdd-v3

Used for high number of spokes, such as for connecting multiple remote sites or mesh topologies.

dictionary-size

Displays the current size of the dictionary, which deduplication uses to look up byte patterns.

[disabled | enabled]

Enables or disables deduplication. The default value is enabled. Note that if you enable or disable deduplication, you must then restart the BIG-IP WOM system using bigstart restart, or the change takes effect the next time the BIG-IP device reboots.

max-endpoint-count

Specifies the maximum number of concurrent remote endpoints supported by symmetric data deduplication. For codec sdd-v3, the system sets the value at 128.

SEE ALSO

sys datastor, list, modify, show, tmsh, wom profile isession,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-09-26 wom deduplication(1)

wom diagnose-conn

NAME

diagnose-conn - Diagnoses network connection problems.

MODULE

wom

SYNTAX

```
run diagnose-conn
```

DESCRIPTION

You can use the `diagnose-conn` component within the `wom` module to display diagnostic information about network connections.

SEE ALSO

`run`, `tmsch`, `wom verify-config`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-04-10 wom diagnose-conn(1)

wom endpoint-discovery

NAME

`endpoint-discovery` - Configures automatic discovery of remote endpoints for WAN optimization.

MODULE

`wom`

SYNTAX

Configure the `endpoint-discovery` component with the `wom` module using the syntax in the following sections.

MODIFY

`modify endpoint-discovery`

options:

`auto-save` [disabled | enabled]
`description` [string]
`discoverable` [disabled | enabled]
`discovered-endpoint` [disabled | enabled]
`icmp-max-requests` [integer]
`icmp-min-backoff` [integer]
`icmp-num-retries` [integer]
`max-endpoint-count` [integer]
`mode` [disable | enable-all | enable-icmp | enable-tcp]

`reset-stats endpoint-discovery`

DISPLAY

`list endpoint-discovery`

`show running-config endpoint-discovery`

options:

`all-properties`
`non-default-properties`
`one-line`

`show endpoint-discovery`

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DESCRIPTION

You can use the `endpoint-discovery` component to specify parameters for automatically discovering remote endpoints for WAN optimization. These endpoints are configured WAN Optimization Managers on remote BIG-IP(r) systems that advertise themselves to the configured WAN Optimization Manager on the local BIG-IP system.

EXAMPLES

`modify endpoint-discovery max-endpoint-count 10`

Limits the number of remote endpoints that can be discovered to ten. After discovering ten remote endpoints, the WOM stops sending probe messages.

`list endpoint-discovery all-properties`

Displays the configuration parameters for the discovery of remote endpoints.

OPTIONS

`auto-save`

Specifies whether the system automatically saves remote endpoints that it discovers. The default value is enabled.

`description`

User defined description.

discoverable
Specifies whether the WAN Optimization Manager responds to probe messages it receives from WAN Optimization Managers on remote BIG-IP systems. The default value is enabled.

discovered-endpoint
Specifies whether the WAN Optimization Manager sends out probe messages to discover other WAN Optimization Managers on remote BIG-IP systems in the network. The default value is enabled.

icmp-max-requests
Specifies the maximum number of ICMP probe message requests, after which the system stops sending probe message requests until at least one message is cleared from the queue by either a timeout or a response. The default value is 1024.

icmp-min-backoff
Specifies the maximum number of seconds to wait before abandoning an ICMP probe message request and resending it. The range is from 0 to 255. The default value is 5.

icmp-num-retries
Specifies the maximum number of times the system sends an ICMP probe message request for a single flow. The range is from 0 to 255. The default value is 10.

max-endpoint-count
Specifies the highest number of endpoints for the system to discover before it stops sending probe messages. The range is from 0 to 255. The default value is 0, which indicates no limit.

mode Specifies the type of probe messages the system should send. The default value is enable-all.

The options are:

disable
Turns off probe messages.

enable-icmp
Sends only ICMP probe messages.

enable-tcp
Sends only TCP probe messages.

enable-all
Sends both ICMP and TCP probe messages.

SEE ALSO

list, modify, show, tmsh, wom local-endpoint, wom remote-endpoint, wom server-discovery

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-04-10 wom endpoint-discovery(1)

wom local-endpoint

NAME
local-endpoint - Configures the local endpoint for the WAN Optimization Manager.

MODULE
wom

SYNTAX
Configure the local-endpoint component within the wom module using the following syntax.

CREATE/MODIFY
create local-endpoint
modify local-endpoint
options:
addresses [add | delete | replace-all-with] {
[ip address]
}
addresses none
allow-nat [disabled | enabled]
description [string]

endpoint [disabled | enabled]
internal-forwarding [disabled | enabled]
ip-encap-mtu [unsigned integer]
ip-encap-profile [none | profile name]
ip-encap-type [gre | ipip | ipsec | none]
no-route [drop | passthru]
server-ssl [none | profile name]
snat [local | none | remote]
tunnel-port [unsigned integer]

DISPLAY

list local-endpoint
show local-endpoint
show running-config local-endpoint
options:
 all-properties
 non-default-properties
 one-line

DELETE

delete local-endpoint

DESCRIPTION

You can use the local-endpoint component to modify the settings for the local endpoint for the WAN Optimization Manager on the local BIG-IP(r) system.

EXAMPLES

modify local-endpoint allow-nat disabled

Disables the allow-nat option, specifying that the system does not accept connections for traffic behind a Network Address Translation (NAT) device.

list local-endpoint all-properties

Displays all of the properties of the local-endpoint component.

OPTIONS

addresses

Specifies a single IP address the system uses for the local endpoint. The IP address must be in the same subnet as a self IP address on the BIG-IP(r) system.

allow-nat

When enabled, specifies that the system accepts connections for traffic behind a Network Address Translation device. The default value is enabled.

description

User defined description.

endpoint

When enabled, specifies that the local endpoint is available for initiating and receiving optimized traffic. The default value is enabled.

To turn off WAN optimization on this endpoint, use disabled.

internal-forwarding

When enabled, specifies that the local endpoint is available for forwarding internal traffic to remote endpoints. The default value is disabled.

This parameter works only if internal-forwarding for remote-endpoint is set to default.

ip-encap-mtu

Specifies the maximum transfer unit for IP encapsulated traffic.

ip-encap-profile

Specifies the name of the profile with the encapsulation settings. This profile must be of the type specified for the setting ip-encap-type.

ip-encap-type

Specifies the type of IP layer encapsulation to perform on iSession(tm) traffic.

The default value is none. The options are:

gre The system uses the Generic Routing Encapsulation (GRE) tunneling protocol.

ipip The system uses the IP over IP (IPIP) tunneling protocol.

ipsec

The system uses IP security (IPsec) encapsulation.

none No IP encapsulation takes place.

no-route

Specifies what the system does with traffic for which there is no remote endpoint to complete the iSession connection.

The default value is passthru. The options are:

drop The system terminates the traffic flow.

passthru

The traffic flow continues without an iSession connection.

server-ssl

Specifies the default server SSL profile the system uses for all encrypted outbound connections. The default value is none.

snat Specifies the IP address the system uses for incoming traffic as the source IP address of the TCP connection between the WAN Optimization Manager and the server.

The default value is none. The options are:

local

The system uses the endpoint IP address closest to the destination. Use this setting to make sure the return route also goes through the BIG-IP system, so that both sides of the connection can be optimized. This setting is useful if responses returning from the server to the client would not normally pass through the BIG-IP system.

none The system uses the original connecting client IP address.

remote

The system uses the source IP address of the incoming iSession connection. Use this setting when an appliance that uses NAT is located between the WAN Optimization endpoints.

tunnel-port

Specifies the number of the port on the local endpoint that the WAN Optimization Manager uses for control connections. The port must have access through the firewall. The range is from 1 to 65535. The default value is 443.

SEE ALSO

list, modify, show, tmsh, wom advertised-route, wom remote-endpoint

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

BIG-IP 2014-05-08 wom local-endpoint(1)

wom profile cifs

NAME

cifs - Configures a Common Internet File System (CIFS) profile.

MODULE

wom profile

SYNTAX

Configure the cifs component within the wom profile module using the syntax shown in the following sections.

CREATE/MODIFY

create cifs [name]

modify cifs [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

fast-close [disabled | enabled]

fast-set-file-info [disabled | enabled]

office-2003-extended [disabled | enabled]

read-ahead [disabled | enabled]

record-replay [disabled | enabled]

write-behind [disabled | enabled]

DISPLAY

list cifs

list cifs [[name] | [glob] | [regex]] ...]

show running-config cifs

show running-config cifs [[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line
partition

show cifs

show cifs [[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DELETE

delete cifs [name]

DESCRIPTION

You can use the cifs component to manage a CIFS profile.

EXAMPLES

create cifs my_cifs_profile

Creates a CIFS profile named my_cifs_profile using the system defaults.

modify cifs my_cifs_profile fast-close disabled

Turns off fast-close for the CIFS profile named my_cifs_profile.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the profile that you want to use as the parent profile. The new profile inherits all settings and values from the parent profile specified. The default value is cifs.

description

User defined description.

fast-close

Specifies whether the system speeds up file close operations by fulfilling them through the WAN Optimization Manager closer to the request initiator. The default value is enabled.

fast-set-file-info

Specifies whether the system speeds up file metadata change requests by fulfilling the requests through the WAN Optimization Manager closer to the request initiator. The default value is enabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

office-2003-extended

Specifies whether the system performs read-ahead operations based on parsing the Microsoft CDF file and understanding its structure. The default value is enabled.

partition

Displays the administrative partition within which the component resides.

read-ahead

Specifies whether the system speeds up CIFS file downloads by prefetching the file data on the WAN Optimization Manager closer to the request initiator. The default value is enabled.

record-replay

Specifies whether the system opens CIFS files faster by performing more intelligent read-ahead operations. The default value is enabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

write-behind

Specifies whether the system speeds up CIFS file uploads to the server by fulfilling write requests through the WAN Optimization Manager closer to the request initiator. The default value is enabled.

SEE ALSO

create, delete, glob, list, ltm virtual, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2013. All rights reserved.

wom profile isession

NAME

isession - Configures an iSession profile.

MODULE

wom profile

SYNTAX

Configure the iSession component within the wom profile module using the following syntax.

CREATE/MODIFY

create isession [name]

modify isession [name]

options:

adaptive-compression [disabled | enabled]

app-service [[string] | none]

compression [disabled | enabled]

compression-codecs [add | delete | none | replace-all-with] {

options:

bzip2

deflate

lzo

}

data-encryption [disabled | enabled]

deduplication [disabled | enabled]

defaults-from [[name] | none]

deflate-compression-level [integer]

description [string]

mode [disabled | enabled]

port-transparency [disabled | enabled]

reuse-connection [disabled | enabled]

target-virtual [none | host-match-all | host-match-no-isession | virtual-match-all]

reset-stats isession

reset-stats isession [[[name] | [blog] | [regex]] ...]

DISPLAY

list isession

list isession [[[name] | [glob] | [regex]] ...]

show running-config isession

show running-config isession [[[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

show isession

show isession [[[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

field-fmt

global

DELETE

delete isession [name]

DESCRIPTION

You can use the isession component to manage an iSession profile.

EXAMPLES

```
create isession my_ession_profile defaults-from isession
```

Creates an iSession profile named my_ession_profile using the system defaults.

```
modify isession my_ession_profile deduplication disabled
```

Turns off deduplication for the iSession profile named my_ession_profile.

OPTIONS

adaptive-compression

Enables or disables the automatic selection of the optimal compression algorithm for the current traffic, based on link speed. The system can use only compression algorithms that are specified. The default value is enabled.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

compression

Enables or disables the compression of data according to the methods you select for the attribute compression-codecs. The default value is enabled.

compression-codecs

Specifies the codecs to use for compression. The following codecs are available:

bzip2

Specifies the use of the bzip2 compression algorithm, which improves compression ratios on low-bandwidth data links.

deflate

Specifies the use of the Deflate data compression algorithm.

lzo Specifies the use of the Lempel-Ziv-Oberhumer (LZO) data compression algorithm.

data-encryption

Enables or disables encryption of the traffic on the outbound connection. If you select enabled, the system uses the SSL profiles specified on the local and remote endpoints of the iSession connection. The default value is disabled.

deduplication

Enables or disables data deduplication, which replaces previously transmitted data with references, thus reducing the amount of bandwidth needed to transfer data over the WAN. The default value is enabled.

defaults-from

Specifies the profile that you want to use as the parent profile. The new profile inherits all settings and values from the parent profile specified. The default value is isession.

deflate-compression-level

Specifies the level of compression, if deflate-compression is specified and adaptive-compression is disabled. The range is 1 to 9. A higher value causes the CPU to spend more time looking for matches, which may result in better compression. The default value is 1.

description

User defined description.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

mode Enables or disables the use of this profile for WAN optimization traffic. The default value is enabled.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

partition

Displays the administrative partition within which the component resides.

port-transparency

Enables or disables the preservation of the destination port specified by the client over the WAN. The default value is enabled.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

reuse-connection

Enables or disables the saving and reuse of connections between the local and remote WAN Optimization Managers. The default value is enabled.

target-virtual

For terminated iSession traffic, specifies the matching criteria that a client-side BIG-IP system uses to select a target virtual server on the server-side BIG-IP system.

The default value is virtual-match-all. The options are:

none Specifies that the system sends the terminated iSession traffic directly to the server.

host-match-all

Specifies that the system selects the closest match from all the host virtual servers.

host-match-no-isession

Specifies that the system matches only host virtual servers with no iSession profile.

virtual-match-all

Specifies that the system selects the closest match from all the virtual servers.

SEE ALSO

create, delete, glob, list, ltm virtual, modify, regex, reset-stats, show, tmsh, wom local-endpoint, wom remote-endpoint

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-05-22 wom profile isession(1)

wom profile mapi

NAME

mapi - Configures a Messaging Application Program Interface (MAPI) profile.

MODULE

wom profile

SYNTAX

Configure the mapi component within the wom profile module using the following syntax.

CREATE/MODIFY

create mapi [name]

modify mapi [name]

options:

app-service [[string] | none]

defaults-from [[name] | none]

description [string]

discover-exchange-servers [disabled | enabled]

native-compression [disabled | enabled]

DISPLAY

list mapi

list mapi [[name] | [glob] | [regex]] ...]

show running-config mapi

show running-config mapi [[name] | [glob] | [regex]] ...]

options:

all-properties

app-service

non-default-properties

one-line

partition

show mapi

show mapi [[name] | [glob] | [regex]] ...]

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DELETE

delete mapi [name]

DESCRIPTION

You can use the mapi component to manage a MAPI profile.

EXAMPLES

create mapi my_mapi_profile

Creates a MAPI profile named my_mapi_profile using the system defaults.

modify mapi my_mapi_profile native-compression enabled

Turns on native-compression for the MAPI profile named my_mapi_profile.

OPTIONS

app-service

Specifies the name of the application service to which the object belongs. The default value is none.

Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

defaults-from

Specifies the profile that you want to use as the parent profile. The new profile inherits all settings and values from the parent profile specified. The default value is mapi.

description

User defined description.

discover-exchange-servers

Enables or disables the automatic discovery of the Microsoft Exchange servers in the network and creation of a virtual server for each one discovered. The default value is disabled.

glob Displays the items that match the glob expression. See help glob for a description of glob expression syntax.

name Specifies a unique name for the component. This option is required for the commands create, delete, and modify.

native-compression

Enables or disables native Microsoft Exchange compression. The default value is disabled.

partition

Displays the administrative partition within which the component resides.

regex

Displays the items that match the regular expression. The regular expression must be preceded by an at sign (@[regular expression]) to indicate that the identifier is a regular expression. See help regex for a description of regular expression syntax.

SEE ALSO

create, delete, glob, list, ltm virtual, modify, regex, show, tmsh

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-05-22 wom profile mapi(1)

wom remote-endpoint

NAME

remote-endpoint - Configures one or more remote endpoints for the WAN Optimization Manager.

MODULE

wom

SYNTAX

Configure the remote-endpoint component within the wom module using the following syntax.

CREATE/MODIFY

create remote-endpoint [name]

modify remote-endpoint [name]

options:

address [ip address]

allow-routing [disabled | enabled]

app-service [[string] | none]

dedup-action [none | cache-refresh]

description [string]

endpoint [disabled | enabled]

internal-forwarding [default | disabled | enabled]

ip-encap-mtu [unsigned integer]

ip-encap-profile [none | profile name]

ip-encap-type [default | gre | ipip | ipsec | none]

origin [configured | discovered | manually-saved | persistable]

server-ssl [none | profile name]

snat [default | local | none | remote]

tunnel-encrypt [disabled | enabled]

tunnel-port [unsigned integer]

reset-stats remote-endpoint

DISPLAY

list remote-endpoint

list remote-endpoint [name]

show running-config remote-endpoint

show running-config remote-endpoint [name]

options:

all-properties

dedup-codec

non-default-properties

one-line

show remote-endpoint
show remote-endpoint [name]
options:
(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

DELETE
delete remote-endpoint [name]

Note: If you delete a remote endpoint without also disabling the endpoint-discovery component, the remote endpoint may reappear as it is rediscovered. To remove a remote endpoint from traffic initiated by this WAN Optimization Manager, set the endpoint option of the remote-endpoint component to disabled.

DESCRIPTION

You can use the remote-endpoint component to create, modify, or delete a remote endpoint for traffic from the local WAN Optimization Manager.

EXAMPLES

modify remote-endpoint 13.16.0.5 endpoint disabled

Disables the WAN optimization connection to the remote endpoint that is named 13.16.0.5.

list remote-endpoint all-properties

Displays all the properties of all the remote endpoints for traffic from the local WAN Optimization Manager.

OPTIONS

allow-routing

Specifies whether there is a route from the local endpoint to this remote endpoint through which the local endpoint can establish connections. The default value is enabled.

address

Specifies the IP address of the remote endpoint.

app-service

Specifies the name of the application service to which the object belongs. The default value is none.
Note: If the strict-updates option is enabled on the application service that owns the object, you cannot modify or delete the object. Only the application service can modify or delete the object.

description

User defined description.

dedup-action

Clears the cache used for symmetric data deduplication on the specified remote endpoint and immediately resets the value to none.

dedup-codec

Displays the deduplication codec used by the remote endpoint: `sdd-v2` or `sdd-v3`.

endpoint

When enabled, specifies that traffic can be optimized between the local and remote endpoints. The default value is enabled.

Note: Disabling a remote endpoint affects only the connection between the local endpoint and this remote endpoint.

internal-forwarding

When enabled, specifies that the remote endpoint is available for forwarding internal traffic.

When disabled, specifies that the remote endpoint is NOT available for forwarding internal traffic.

When default, specifies that forwarding internal traffic is managed by the `local-endpoint.internal-forwarding` setting.

The default value is default.

ip-encap-mtu

Specifies the maximum transfer unit for IP encapsulated traffic. The default value is 0.

ip-encap-profile

Specifies the name of a profile with encapsulation settings. This profile must be of the type specified for the setting `ip-encap-type`.

ip-encap-type

Specifies the type of IP layer encapsulation performed on iSession traffic.

The default value is default. The options are:

default

The system uses the `ip-encap-type` value set for the local endpoint.

gre The system uses the Generic Routing Encapsulation (GRE) tunneling protocol.

ipip The system uses the IP over IP (IPIP) tunneling protocol.

ipsec

The system uses IP security (IPsec) encapsulation.

none No IP encapsulation takes place.

origin
Specifies whether the remote endpoint was discovered automatically or configured manually.

The options are:

configured
Indicates that you manually configured this remote endpoint. The system automatically sets this value, and you cannot change it.

discovered
Indicates that the system automatically discovered this remote endpoint. Note that endpoints for which the value of the origin property is discovered are not saved to the file `bigip_local.conf`.

manually-saved
After you run the command `save / sys config`, the value of the origin property that was set to persistable changes to manually-saved. Note that after the system changes the value to manually-saved, you cannot change it again.

persistable
Change the origin from discovered to persistable, if you want to save the endpoint to the file `bigip_local.conf` when you use the command `save / sys config`.

server-ssl
Specifies the server SSL profile the system uses to connect to this remote endpoint. This setting overrides the server-ssl setting for the local-endpoint component. The default value is none.

snat Specifies the IP address the system uses as the source IP address of the TCP connection between the WAN Optimization Manager and the server.

The default value is default. The options are:

default
The system uses the snat value set for the local-endpoint component.

local
The system uses the endpoint IP address closest to the destination. Use this setting to make sure the return route also goes through the BIG-IP system, so that both sides of the connection can be optimized. This setting is useful if responses returning from the server to the client would not normally pass through the BIG-IP system.

none The system uses the original connecting client IP address.

remote
The system uses the source IP address of the incoming iSession connection. Use this setting when an appliance that uses NAT is located between the WAN Optimization Manager endpoints.

tunnel-encrypt
Enables or disables encryption of traffic passing between the two WAN Optimization Managers. The default value is enabled

tunnel-port
Specifies whether to use a specific port for traffic optimized to this endpoint or to use port transparency (0). The default value is 443.

SEE ALSO

`create`, `delete`, `list`, `modify`, `show`, `tmsh`, `wom advertised-route`, `wom local-endpoint`

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2014-05-08 wom remote-endpoint(1)

wom remote-route

NAME

`remote-route` - Displays the destination routes learned from the remote endpoints.

MODULE

wom

SYNTAX

Display the remote-route component within the wom module using the syntax in the following section.

DISPLAY

show remote-route

options:

(default | exa | gig | kil | meg | peta | raw | tera | yotta | zetta)

detail

DESCRIPTION

You can use the remote-route component to view the subnets that the system can reach through the remote endpoint(s). The system can optimize traffic destined for these subnets.

EXAMPLES

show remote-route

Displays the subnets reachable through the remote endpoint(s) configured on the WAN Optimization Manager.

show remote-route detail

Displays detailed information about the remote endpoint(s) through which the displayed subnets can be reached.

SEE ALSO

show, tmsh, wom advertised-route, wom remote-endpoint, wom server-discovery,

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-04-11 wom remote-route(1)

wom server-discovery

NAME

server-discovery - Configures the dynamic discovery of servers that can be reached through the local endpoint and the routes to reach them.

MODULE

wom

SYNTAX

Configure the server-discovery component within the wom module using the syntax in the following sections.

MODIFY

modify server-discovery

options:

auto-save [disabled | enabled]

description [string]

filter-mode [exclude | include]

idle-time-limit [integer]

ip-ttl-limit [integer]

max-server-count [integer]

min-idle-time [integer]

min-prefix-length-ipv4 [integer]

min-prefix-length-ipv6 [integer]

mode [disabled | enabled]

rtt-threshold [integer]

subnet-filter [add | delete | none | replace-all-with] {

[ip address]

}

time-unit [days | hours | minutes]

DISPLAY

list server-discovery

show running-config server-discovery

options:

all-properties

auto-save

current-module

description

filter-mode

idle-time-limit

ip-ttl-limit
max-server-count
min-idle-time
min-prefix-length-ipv4
min-prefix-length-ipv6
mode
non-default-properties
one-line
rtt-threshold
subnet-filter
time-unit

DESCRIPTION

You can use the server-discovery component to configure the dynamic discovery of servers and the routes to reach them through the local endpoint. The local endpoint advertises these routes to any remote endpoints to which it is connected.

EXAMPLES

list server-discovery all-properties

Displays the settings for dynamic discovery of advertised routes.

modify server-discovery mode disabled

Disables the dynamic discovery of advertised routes.

OPTIONS

auto-save

Specifies whether the system automatically saves the subnets that it discovers that can be reached through the local endpoint. The default value is enabled.

description

User defined description.

filter-mode

Specifies whether the subnets you add using the attribute **subnet-filter** are excluded from or included in the discovery of advertised routes. If you specify **include**, and do not specify any IP addresses, no subnets are discovered. The default is **exclude** with no IP addresses specified, which means that all advertised routes that conform to the specified attributes are discovered.

idle-time-limit

Specifies the maximum length of time a route can be idle without being removed from discovery. The default value is 0. Use the attribute **time-unit** to set the unit of measure. Use the attribute **min-idle-time** to set the minimum length of idle time.

ip-ttl-limit

Specifies the number of network segments on which a packet is allowed to travel before the route is removed from discovery. The more routers a packet travels through, the smaller the ip ttl value is. The range is 0 to 255. The default value is 5.

max-server-count

Specifies the highest number of servers the system discovers before it stops looking. The default value is 50.

min-idle-time

Specifies the minimum length of time a route must be idle before being removed from discovery. The default value is 0, which indicates that idle time is not considered in discovery. Use the attribute **time-unit** to set the unit of measure. Use the attribute **idle-time-limit** to set the maximum length of idle time.

min-prefix-length-ipv4

Specifies the minimum prefix length for route aggregation in IPV4 networks. The range is 0 to 32. The default value is 32.

min-prefix-length-ipv6

Specifies the minimum prefix length for route aggregation in IPV6 networks. The range is 0 to 128. The default value is 128.

mode Enables or disables the dynamic discovery of servers that can be reached through the local endpoint. For server discovery to take place, the setting **mode** of the component **wom endpoint-discovery** must not be set to **disabled**.

rtt-threshold

Specifies that the system does not add servers it discovers with a round-trip time greater than this value, in milliseconds. The default value is 10.

subnet-filter

Specifies the IP addresses of the subnets to include in or exclude from the discovery of advertised routes, depending on the setting you selected for the attribute **filter-mode**. The default is **none**. If you selected **include** for the attribute **filter-mode**, and do not specify any IP addresses, no subnets are discovered.

time-unit

Specifies the unit of measure (days, hours, or minutes) for the length of idle time specified using the attributes **idle-time-limit** and **min-idle-time**.

SEE ALSO

list, modify, show, tmsh, wom advertised-route, wom endpoint-discovery, wom local-endpoint, wom remote-route

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2010-2012. All rights reserved.

BIG-IP 2012-04-11 wom server-discovery(1)

wom verify-config

NAME

verify-config - Checks the WAN Optimization Manager configuration.

MODULE

wom

SYNTAX

run verify-config

DESCRIPTION

You can use the verify-config component within the wom module to display configuration information about the WAN Optimization Manager that can be used for troubleshooting.

SEE ALSO

run, tmsh, wom diagnose-conn

COPYRIGHT

No part of this program may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the express written permission of F5 Networks, Inc.

F5 Networks and BIG-IP (c) Copyright 2009-2012. All rights reserved.

BIG-IP 2012-04-11 wom verify-config(1)
